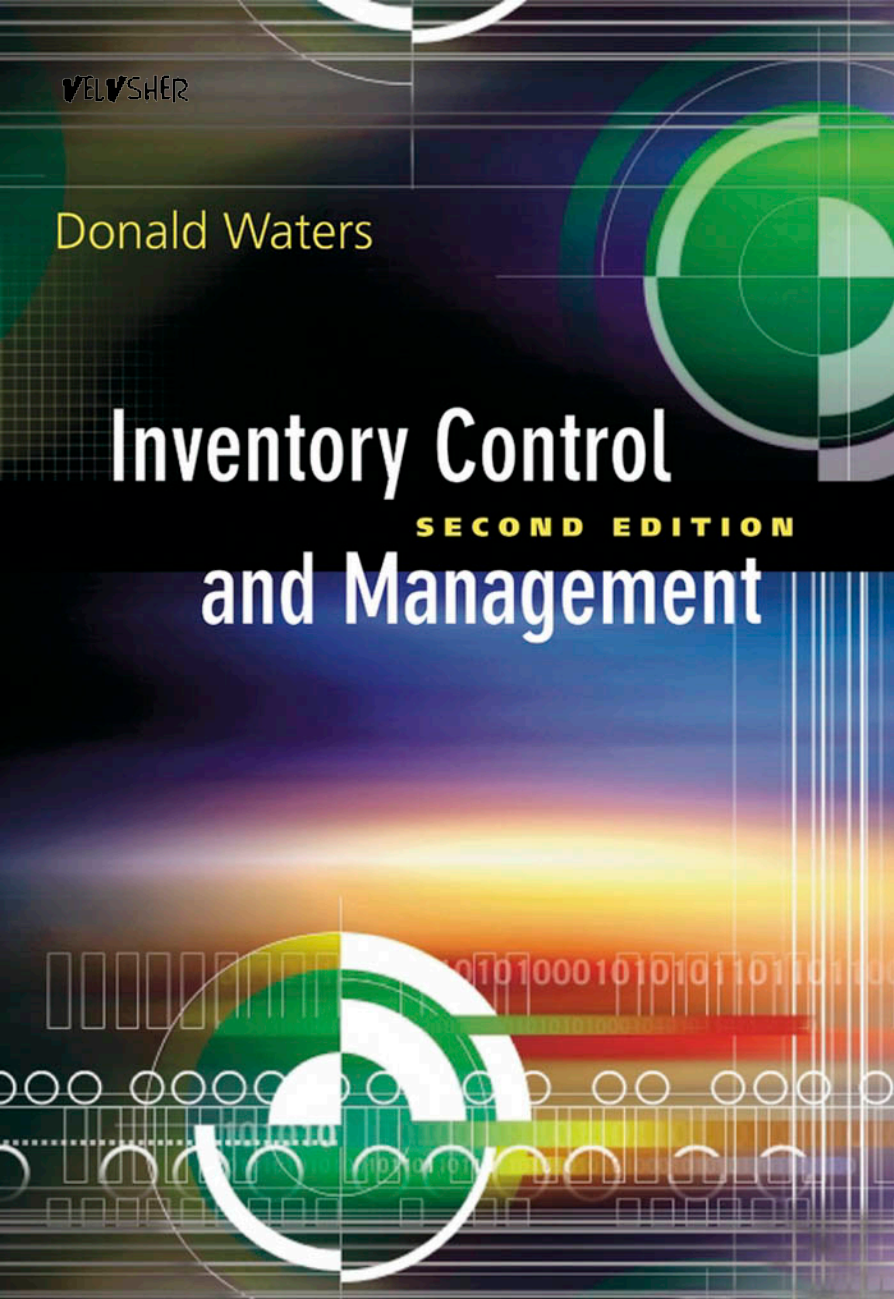

VELVSHER
Donald Waters
Inventory Control
SECOND EDITION
and Management

Inventory Control and Management

Donald Waters



Inventory Control and Management

Inventory Control and Management

Donald Waters



Copyright © 2003 John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester,
West Sussex PO19 8SQ, England

Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk
Visit our Home Page on www.wileyeurope.com or www.wiley.com

All Rights Reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP, UK, without the permission in writing of the Publisher. Requests to the Publisher should be addressed to the Permissions Department, John Wiley
& Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770620.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 33 Park Road, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809

John Wiley & Sons Canada Ltd, 22 Worcester Road, Etobicoke, Ontario, Canada M9W 1L1

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Cataloging-in-Publication Data

Waters, C. D. J. (C. Donald J.), 1949–
Inventory control and management / by Donald Waters. – 2nd ed.
p. cm.
Includes bibliographical references and index.
ISBN 0-470-85876-1 (alk. paper)
1. Inventory control. I. Title.

TS160.W38 2003
658.7′87–dc21

2003057186

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN 0-470-85876-1

Typeset in 10/12pt Palatino by Laserwords Private Limited, Chennai, India
Printed and bound in Great Britain by TJ International, Padstow, Cornwall
This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

To Jack

Contents

Preface . . . . . **xi**

PART I

Introduction **1**

1 Stocks and Inventories **3**

Aims of the chapter 3
Stocks of materials 3
Reasons for holding stocks 7
Stocks in the supply chain 11
Trends affecting stock 19
Changes to aggregate stocks 24
Chapter review 28
Project 29
Discussion questions 29
References and further reading 29

2 Stocks within an Organization **31**

Aims of the chapter 31
Inventory management and logistics 31
Setting the aims of inventory management 36
Strategic role of stock 44
Costs of holding stock 48
Approaches to inventory control 55
Chapter review 58
Project 59
Discussion questions 60
References and further reading 60

PART II
Methods for Independent Demand 63

3 Economic Order Quantity 65

Aims of the chapter 65
Defining the economic order quantity 65
Adjusting the economic order quantity 77
Uncertainty in demand and costs 86
Adding a finite lead time 89
Chapter review 96
Project 97
Problems 97
Discussion questions 98
References and further reading 98

4 Models for Known Demand 99

Aims of the chapter 99
Price discounts from suppliers 100
Finite replenishment rate 113
Planned shortages with back-orders 120
Lost sales 125
Constraints on stock 128
Discrete, variable demand 135
Chapter review 141
Project 142
Problems 142
Discussion questions 144
References and further reading 144

5 Models for Uncertain Demand 147

Aims of the chapter 147
Uncertainty in stocks 148
Models for discrete demand 154
Order quantity with shortages 166
Service level 170
Uncertain lead time demand 173
Periodic review methods 181
Chapter review 187
Project 188
Problems 188
Discussion questions 189
References and further reading 190

PART III
Information for Inventory Management 193

6 Sources of Information 195

Aims of the chapter 195
Inventory management information systems 195
Information from accounting 202
Information about supply and demand 211
Warehousing 218
Chapter review 224
Project 225
Problems 225
Discussion questions 226
References and further reading 227

7 Forecasting Demand 229

Aims of the chapter 229
Methods of forecasting 230
Judgemental forecasts 234
Time series 237
Causal forecasting 241
Projective forecasting 248
Planning forecasts 260
Chapter review 262
Project 263
Problems 264
Discussion questions 265
References and further reading 266

8 Planning and Stocks 267

Aims of the chapter 267
Levels of planning 268
Aggregate planning 274
Master schedules 285
Operational schedules 288
Simulation of stocks 292
Chapter review 299
Project 299
Problems 300
Discussion questions 302
References and further reading 302

PART IV
Methods for Dependent Demand **305**

9 Material Requirements Planning **307**

Aims of the chapter 307
Limitations of independent demand methods 307
Approach of material requirements planning 310
Benefits and problems with MRP 321
Adjusting the MRP schedules 327
Extensions to MRP 332
Chapter review 336
Project 337
Problems 337
Discussion questions 339
References and further reading 339

10 Just-in-Time **341**

Aims of the chapter 341
Principles of just-in-time 341
Main features for stocks 345
Achieving just-in-time operations 348
Other effects of JIT 354
Benefits and disadvantages of JIT 363
Extending JIT along the supply chain 365
Comparisons with other methods of inventory management 367
Chapter review 370
Project 371
Problems 371
Discussion questions 372
References and further reading 373

Index **385**

Preface

The subject

This is a book about inventory management. It describes recent thinking about stock and methods for its control.

You might imagine stock as warehouses full of goods but every organization holds stock, even those providing the most intangible service. A small company of knowledge workers, for example, stores information and experience, and it faces the same problems of inventory management as a giant manufacturer with its stores of finished goods and components.

As stocks are almost universal, we should start with the basic question, 'Why do organizations hold stocks?' The main answer is to allow for variations and uncertainty in supply and demand – they give a buffer between suppliers and customers, maintaining customer service even when there are problems in the supply chain. Unfortunately, this safeguard comes at a high price, and organizations are continually looking for ways of reducing their inventory costs without affecting service. In recent years this search has led to many changes. 'Scientific inventory control' is still at the core of inventory management, but it has been enhanced by requirements planning and just-in-time, with e-commerce giving a fast and efficient flow of material through an integrated supply chain.

Stock does not exist in isolation, so we have to consider its impact on other parts of the organization. There are no clear lines between inventory management and, say, procurement, supply chain management, warehousing, or broader operations. So we have to set inventory management in its overall context, noting its interactions with other activities, and explicitly recognizing its strategic importance.

Obviously we have to put some limit on the material covered, and we have concentrated on the core questions of inventory management: What exactly is inventory management? How do decisions about stock affect other operations? How can we control stocks? What information do we need? What is the effect of new methods and technology? The answers to these questions embrace the most important issues of inventory management.

About the book

The book gives a self-contained introduction to inventory management. It develops the subject from basic principles through to advanced material and newer

developments. It does not assume any previous knowledge of the subject, and anyone reading the book will find it easy to build a detailed understanding of the topic.

The book is also careful not to assume any specific knowledge about management, operations, mathematics, accounting, or any other subject. This means that many different types of people can usefully read it, including students doing a general course in management, business, commerce, or related field. It can also be used for more specialized courses in, say, marketing, supply chain management, operations management, operational research, management science, or production. In addition, it is useful for practising managers who want to learn more about inventory management and how the ideas can be used in their work. Whatever your background, you can use this accessible and user-friendly book to learn about current thinking and practices in inventory management.

The book includes the following features:

- It is an introductory text and assumes no previous knowledge of inventory management, or experience in the area.
- It is clearly written and takes a user-friendly approach, discussing ideas in an informative and easy style.
- It develops material in a logical order, introducing ideas through examples and avoiding abstract discussion and unnecessary mathematics.
- It includes many features, such as chapter aims, key concepts, worked examples, review questions and solutions, projects, discussion questions, additional problems and sources of further reading.
- It is suitable for many types of people, including students doing a wide range of courses and practising managers.
- Ideas about inventory management are changing very quickly, and this book gives an up-to-date view of thinking and practices.
- It takes a broad view of the subject, discussing the context and links with other functions, including the strategic role of stock.
- It includes associated material in an instructors' manual, with a glossary, solutions to problems, spreadsheets, comments on discussion questions, etc.

Overview of the book

The book follows a logical path through inventory management. To make this easier we have divided the material into four parts. Part I gives an overall introduction to inventory management. It discusses the movement of materials through supply chains, reasons for holding stocks, their importance, use, costs, and so on. There are essentially two approaches to inventory management:

- independent demand methods, which assume that overall demand for an item is made up of a large number of independent, small demands from individual customers;
- dependent demand methods, which assume that there are some links between the demands for different items.

Part II of the book looks at independent demand inventory methods. This starts with the classic models of inventory control, which calculate an optimal order size under various conditions. Then Part III discusses the information needed to support these methods, including information from the inventory management information system, forecasts of demand and planned operations. Part IV describes dependent demand methods. The main approaches of this type are based on material requirements planning and just-in-time operations. Together, these four parts give a comprehensive review of current thinking and practices in inventory management.

There are ten chapters in the book, each of which follows a consistent pattern with:

- a list of aims, showing what you should be able to do after reading the chapter;
- a statement of the key concepts covered in the chapter;
- the main material in the chapter, divided into coherent sections;
- summaries and review questions at the end of each section, with solutions to review questions at the end of the book;
- worked examples to illustrate quantitative ideas;
- a review at the end of each chapter summarizing the material covered;
- a short project to encourage research into the subject;
- problems to give practice in the quantitative analyses;
- discussion questions to stimulate thinking and research.

At the end of the book is a set of references and selected further readings.

As well as the main text there is associated material which includes spreadsheets used for calculations, solutions to problems, comments on projects and discussion questions, slide masters, a glossary, review of equations, etc. Details of this are given at the publisher's website *www.wiley.co.uk/waters* or you can contact the author at *donaldwaters@lineone.net*. If you have any comments, queries, requests or suggestions for the book or associated material, the author and publisher would be very pleased to hear them.

Changes to this edition

This edition contains many changes since the first edition was published in 1992. The whole book has been rewritten to make it clearer, add new topics,

reflect changing importance, and generally update material. Some specific changes include:

- complete rewriting to improve the flow and make parts of the text even clearer;
- adding new material that has become important, such as the strategic impact of stocks, integration along the supply chain, e-business, information flows, lean operations, etc.;
- removal of material that has become dated and less relevant;
- use of software – particularly spreadsheets – for all calculations and the Internet as a source of information for discussions and exercises;
- general updating of material, examples and illustrations;
- more features, including reviews, projects, key concepts, supplements, etc.;
- improved design giving the book a better, clearer layout;
- materials for students and instructors on an associated website.

Part I

Introduction

1

Stocks and Inventories

Aims of the chapter

In this chapter we introduce the ideas that lie behind inventory management. We define the terms used, describe the general features of stocks, their purpose, importance and use. We describe some changes that have affected inventory management in recent years. Our aim is to set the scene for later chapters, taking a broad overview before moving on to more detailed discussions.

After reading this chapter you should be able to do the following:

- define the main terms used for inventory management;
- describe the importance of stocks in an organization;
- discuss the reasons for holding stock;
- review the role of stocks in a supply chain;
- explain the benefits of co-ordinated stocks in a supply chain;
- describe some important business trends that affect stocks;
- say how views of stock have changed over time;
- describe the changing pattern of stocks at a national level.

This chapter discusses the following concepts:

- *Stocks*, which are stores of materials that are kept until needed.
- *Inventory*, which is a list of items held in stock.
- *Inventory management*, which is responsible for all aspects of stock control.

Stocks of materials

Definition of terms

All organizations hold stocks. These are the stores of materials they keep until needed. A shop, for example, buys goods from a wholesaler and keeps them in

stock until it sells them to customers; a factory keeps a stock of raw materials for its products; a television company has a stock of recorded programmes; a farmer stores hay to feed his animals over the winter; a research company has a stock of information; a bank holds cash for its day-to-day transactions. Whenever an organization has materials that it does not use immediately, it puts them into stock.

- *Stock* consists of all the goods and materials that are stored by an organization. It is a store of items that is kept for future use.
- An *inventory* is a list of the items held in stock.

An immediate problem is that people use these terms in different ways. In recent years it has become more common to use 'inventory' for both the list of items and the stock itself, and the two terms then become interchangeable. At the same time, organizations refer to their stock as stores, provisions, stockpiles, holdings, reserves, accumulated materials, banks, or a host of other names. To add to the confusion some groups put slightly different interpretations on the terms. Accountants, for example, view 'inventory' as the amount of money tied up in stocks, rather than the stocks themselves, or it might be the total value of an organization's assets. To finance people, 'stocks' are a way of raising capital – in the sense of 'stocks and shares' – and have nothing to do with stores of materials. Usually, these differences are fairly obvious and cause few problems, but sometimes you have to be a bit more careful. In this book we will stick to the standard definitions, where an inventory is a list of the items held in stock.

Each entry in the inventory is a distinct *item* that is held in stock. A supermarket, for example, has 'one-litre bottles of Diet Coke' as a distinct item. Other items in its inventory might be 'two-litre bottles of Diet Coke', 'half-litre bottles of Diet Coke', 'one-litre bottles of Diet Pepsi', and every other distinct product that it sells. A typical supermarket stocks about 30,000 items. Again, some people use different terms, with the most common alternative being *stock keeping unit* or *SKU*.

Each item is sold in standard quantities, or *units*. With our one-litre bottles of Diet Coke, the unit is clearly a bottle. Similarly, '500-gramme tins of Heinz baked beans' is an item in a supermarket, and each tin of beans is a unit; £1 stamps are an item in a post office, and each stamp is a unit; unleaded petrol is an item in a filling station, and each litre is a unit.

- An *item* is a distinct product that is kept in stock: it is one entry in the inventory.
- A *unit* is the standard size or quantity of an item.

Stocks are fairly obvious when you see a shop full of goods, or a warehouse of finished goods. These stocks are tangible and readily identifiable. Sometimes the stocks are a little less obvious, such as the reserves of cash held by a bank,

reservoirs used by a water company, substitutes available for a football team, or seats available in a theatre. An even broader view includes services with intangible stocks, such as the information that is held by research companies, the stock of expertise with consultants, and the store of knowledge in universities. In principle, all of these stocks need the same kinds of management. It is easier to imagine stocks of tangible goods, but remember that in different circumstances organizations can hold stocks of raw materials, components, finished products, people, information, paperwork, messages, knowledge, consumables, energy, money and anything else they need. For simplicity, we will use the general term 'material' for anything that is kept in stock. In the same way, we will always refer to 'organizations' holding stock to cover all types of company, whether it is a not-for-profit organization, a government body, a charity, a quango, a club, or any other body.

Stock cycles

Stocks are formed whenever an organization acquires materials that it does not use immediately. A common practice has a delivery of material arriving from a supplier, and this is kept in stock until needed. Sometimes it is easier to picture a specific operation, so you might imagine the stocks in a supermarket. Goods are delivered by lorry at night, these are checked, sorted and put onto shelves. Then they stay on the shelves until customers buy them. At some point stocks get low, and the supermarket arranges another delivery (as shown in Figure 1.1).

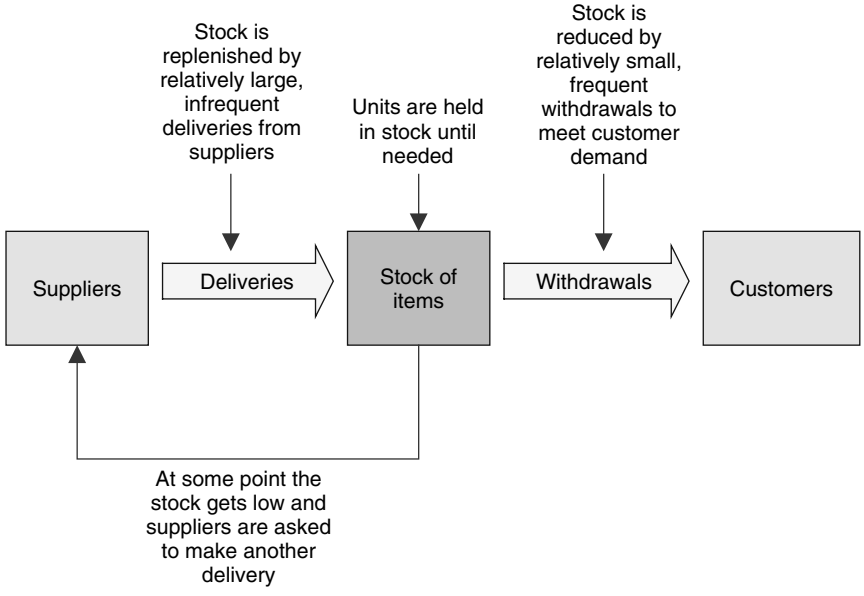

Stock is replenished by relatively large, infrequent deliveries from suppliers
Units are held in stock until needed
Stock is reduced by relatively small, frequent withdrawals to meet customer demand
Suppliers
Deliveries
Stock of items
Withdrawals
Customers
At some point the stock gets low and suppliers are asked to make another delivery

Figure 1.1 A typical use of stock

This sequence of stock replenishment and reduction to meet demand is repeated continuously in a *stock cycle*. Typically, each cycle has the following elements:

1. An organization buys a number of units of an item from a supplier.
2. At an arranged time, these units are delivered.
3. Unless they are needed immediately, the units are put into storage, replenishing the stock.
4. Customers, either internal or external, create demands for the item.
5. Units are removed from stock to meet these demands.
6. At some point, the stock gets low and it is time for the organization to place another order.

Usually deliveries from suppliers are relatively large and infrequent, while demands from customers are smaller and more numerous, giving the typical

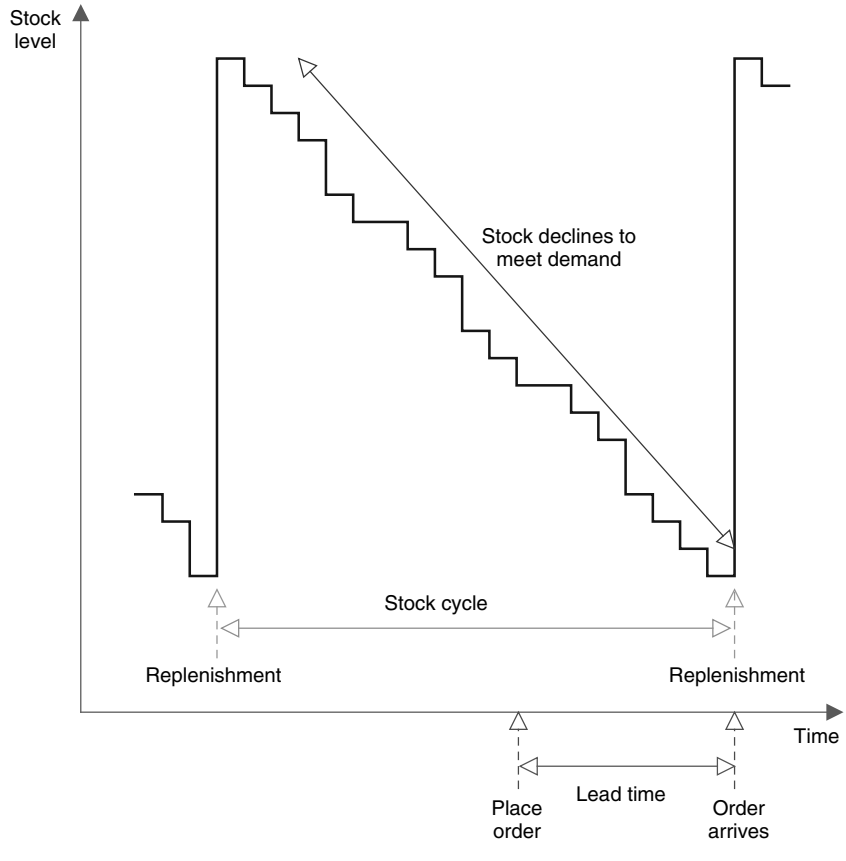


Figure 1.2 Stock levels in a typical cycle

pattern shown in Figure 1.2. Remember that in this broad sense, a *customer* is anyone or anything whose demand is met by removing units from stock. The customer can be internal, when they are someone else within the same organization, or external, when they come from outside the organization. A *supplier* is anyone or anything that replenishes or adds to stock, and again it can be either internal or external. The length of a stock cycle can vary between a few hours (like newspapers and milk which have frequent deliveries) and decades (like gold in Fort Knox that is rarely passed on to customers).

As materials move through the stock cycle, there are many different arrangements for purchasing, storage and delivery. One common feature, however, is that holding stock is surprisingly expensive. We will look at costs in the next chapter, but a rule of thumb says that the cost of holding stock is about 20 per cent of its value a year. If you keep £500 of food in a freezer, it costs about £100 a year. To appreciate the scale of these costs to industry, you only have to look at a big logistics centre, watch convoys of delivery lorries moving along motorways, or realize that Tesco keeps a billion pounds of stock in its shops. Not surprisingly, organizations put a lot of effort into controlling these costs through careful *inventory management*. This function is also called *stock control* or *inventory control*.

- *Inventory management* is the function responsible for all decisions about stock in an organization.
- It makes decisions for policies, activities and procedures to make sure the right amount of each item is held in stock at any time.

Summary

Every organization holds stocks of materials. These are the stores of items – listed in an inventory – that are kept until needed. Stocks are replenished by deliveries from suppliers and reduced to meet demands from customers. Inventory management is responsible for all aspects of stock control.

Review questions

1.1 What is the difference between stock and inventory?

1.2 If a stock item is 'Everyperson's Encyclopaedia', what is a unit?

1.3 How would you define 'suppliers' and 'customers'?

Reasons for holding stocks

Giving a buffer

Stocks are expensive, because of the costs of tied-up capital, warehousing, protection, deterioration, loss, insurance, packaging, administration, and so on. An obvious question, then, is, 'Why do organizations hold stock?' There are several

answers to this, all based on the need for a buffer, or cushion, between supply and demand.

We can illustrate the need for a buffer by considering the stock of bread at Angela's Bakery Shop. It takes some time to make bread, but customers will not wait this time and want a loaf available as soon as they enter the shop. Angela clearly has to plan her baking in advance. If she knows exactly when customers want bread, she can schedule the baking so that loaves are ready at exactly the right time. This would have the advantages of eliminating stock, giving customers the freshest possible bread, and having no leftovers to go stale. In practice, of course, she does not know exactly when customers will buy bread, or how much they want. There is always some variation and uncertainty in the timing and size of customer purchases, and to allow for this Angela bakes loaves in advance and keep a stock on her shelves. Another important concern is that each customer only buys a small amount, but the most efficient way of making bread is in batches of an oven-full at a time. The stock allows for this mismatch between the best rate of supply and actual demand.

Now consider another example with two consecutive operations on an assembly line. Ideally, the first operation finishes a unit, and passes it to the second operation, which starts work immediately. But if the first operation develops a fault, or there is something wrong with the unit, or there is some other reason for a delay in passing the unit forward, the second operation has nothing to work on and it sits idly waiting. The way to avoid this loss of production is to have a small stock of work in progress between the operations. When there are problems with moving a unit forward, the second operation continues working on this stock, and the buffer 'decouples' their operations.

These two examples show how stock gives a buffer between supply and demand. It allows for variation and uncertainty in both supply and demand, and lets operations continue smoothly when there are problems (see Figure 1.3). We can add some details to this idea of a buffer and say that organizations hold stocks to do the following:

- allow for demands that are larger than expected, or at unexpected times;
- allow for deliveries that are delayed or too small;
- allow for mismatches between the best rate of supply and actual rate of demand;
- decouple adjacent operations;
- avoid delays in passing products to customers;
- take advantage of price discounts on large orders;
- allow the purchase of items when the price is low and expected to rise;
- allow the purchase of items that are going out of production or are difficult to find;
- make full loads for delivery and reduce transport costs;
- give cover for emergencies.

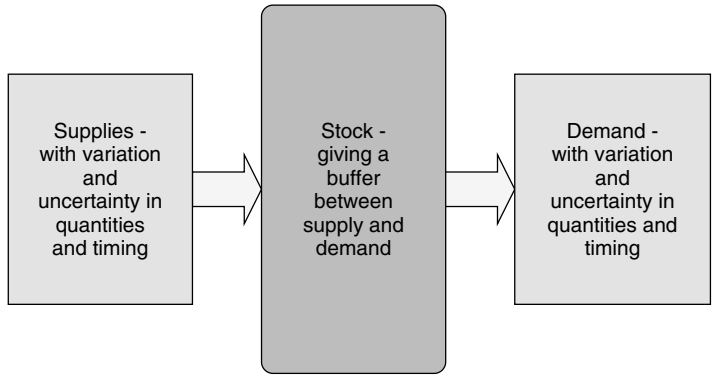

Supplies - with variation and uncertainty in quantities and timing
Stock - giving a buffer between supply and demand
Demand - with variation and uncertainty in quantities and timing

Figure 1.3 Stock gives a buffer between supply and demand

Types of stocks

To achieve these purposes, organizations hold different types of stock. A useful classification has:

- *raw materials*, which have arrived from suppliers and are kept until needed for operations;
- *work in progress*, which are units currently being worked on;
- *finished goods*, which are waiting to be shipped to customers.

This is a fairly arbitrary classification, as one organization's finished goods are another organization's raw materials. Some organizations (notably retailers and wholesalers) have stocks of finished goods only, while others (manufacturers, say) have all three types. Some stock items do not fall easily into these categories, so we can define two additional types as:

- *spare parts*, for machinery, equipment, etc.,
- *consumables*, such as oil, paper, cleaners, etc.

These are needed to support operations, but they do not form a part of the final product (shown in Figure 1.4).

To take a specific example, GlenMorray Knitwear make a range of golf clothes, and their raw materials are wool, cotton, fabrics and other materials waiting to be made into articles; work in progress is the articles being worked on at the moment; finished goods are articles waiting to be delivered to customers; spare parts are kept for the knitting machines and other equipment; and consumables include cleaners, stationery and other material to keep the operations going.

Another less widely used classification of stock describes its overall purpose:

- *Cycle stock* is the normal stock used during operations.
- *Safety stock* is a reserve of materials that is held for emergencies.

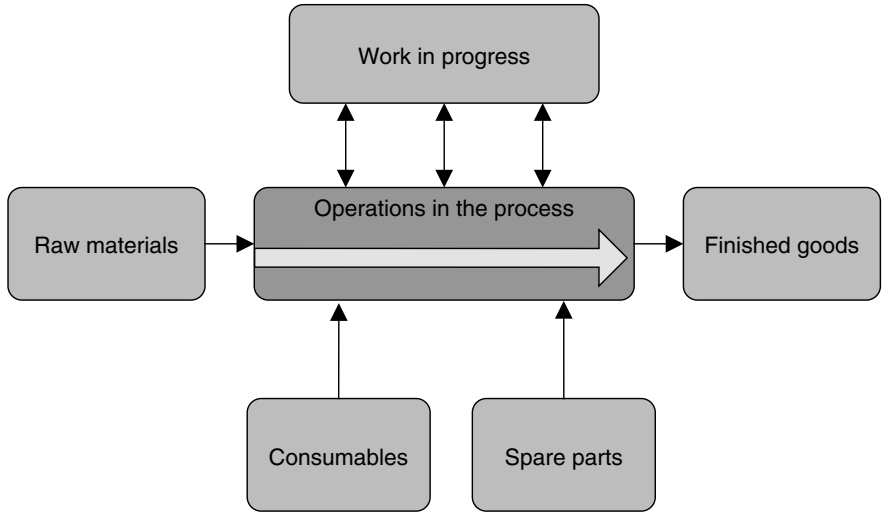


Figure 1.4 Different types of stock

- *Seasonal stock* is used to maintain stable operations through seasonal variations in demand.
- *Pipeline stock* is currently being moved from one location to another.
- *Other stock* consists of all the stocks that are held for some other reason.

Importance of stock

There is a huge variation in the stockholdings of different industries and organizations. Building materials, such as sand and gravel, need fairly large storage areas, but virtually no special attention; expensive items, such as gold and diamonds, need small storage areas, but with high security; perishable goods, such as frozen foods, need special types of storage; information can be stored in huge quantities, but it must allow rapid searching, sorting and retrieval. Despite these differences, you can see that stocks play an important – and even essential – role in every organization. Without stocks most operations are simply impossible. At the very least, stocks allow operations to become more efficient and productive. Stocks affect lead times and availability of materials – thereby affecting customer service, satisfaction, and the perceived value of products. They affect operating costs – and hence profit, return on assets, return on investment and just about every other measure of financial performance. They affect broader operations, by determining the best size, location and type of facilities; they can be risky, because of storage requirements, safety, health and environmental concerns; they can encourage growth of other organizations, such as suppliers and intermediaries offering specialized services

To put it simply, without stocks, organizations could not work. The important question, then, is not whether to hold stocks, but how to manage the stocks that must be held.

Summary

The main purpose of stock is to give a buffer between supply and demand. This safety cushion is essential to ensure the smooth running of operations. Stocks can be raw materials, work in progress, finished goods, spare parts or consumables. The amounts held have widespread effects on the performance of an organization.

Review questions

1.4 How do stocks act as a buffer between operations?

1.5 If suppliers were reliable, there would be no need for stock. Do you think this is true?

1.6 How would you classify lubricating oil for an engine?

Stocks in the supply chain

Shape of supply chains

We have talked about stocks in a single organization, but no organization works in isolation. Each becomes a customer (when buying materials from suppliers) and a supplier (when delivering materials to customers). A wholesaler, for example, acts as a customer when buying goods from manufacturers, and then as a supplier when selling goods to retail shops. Products move through a series of organizations and operations as they travel between original suppliers and final customers. Milk moves through a farm, tanker collection, dairy, bottling plant, distributor and supermarket before we buy it. A toothbrush starts its journey with a company extracting crude oil, and then it passes through pipelines, refineries, chemical works, plastics companies, manufacturers, importers, wholesalers and retailers before finishing up in your bathroom. This series of activities and organizations forms the product's *supply chain*. The function that has overall responsibility for moving materials through the supply chain is *logistics* or *supply chain management*.

- A *supply chain* consists of the series of activities and organizations that materials move through on their journey from initial suppliers to final customers.
- *Logistics* or *supply chain management* is the function responsible for this flow of materials.

As logistics has overall responsibility for the movement – and storage – of materials, inventory management becomes one of the tasks of this broader function.

It is certainly impossible to separate inventory management from other decisions about the supply chain. When we talk about controlling the stock of, say, raw materials, we have to consider the transport of materials, warehousing, purchasing and other activities of supply chain management.

Every product has its own unique supply chain, with materials moving through raw materials suppliers, manufacturers, finishing operations, logistics centres, warehouses, third party operators, transport companies, wholesalers, retailers and a whole range of other operations. In a simplified view, the supply chain for a product consists of tiers of suppliers feeding materials from original sources into its operations, and then tiers of customers moving materials out to the final customers (as shown in Figure 1.5).

There are many variations on this basic model, but the two main features are the supply chain's length and breadth. Here the length refers to the number of tiers, or intermediaries, that materials flow through between source and destination. When farmers sell their produce directly to final customers there is a very short supply chain; on the other hand, computers combine parts from around the world and have long chains. Supply chain breadth is the number of parallel routes that materials can move through on their way to final customers. Cadbury's has a broad supply chain, which means you can buy their chocolate in a huge number of retailers; Pigalle et Fils has a very narrow chain and they only sell their chocolate in two shops in Belgium.

A key point is that every organization on a supply chain holds its own stocks. If the supply chain is very long, or very broad, there is a lot of material held in storage and this is likely to move slowly towards final customers. So one factor in the design of a supply chain is the total amount of stock held. An empirical

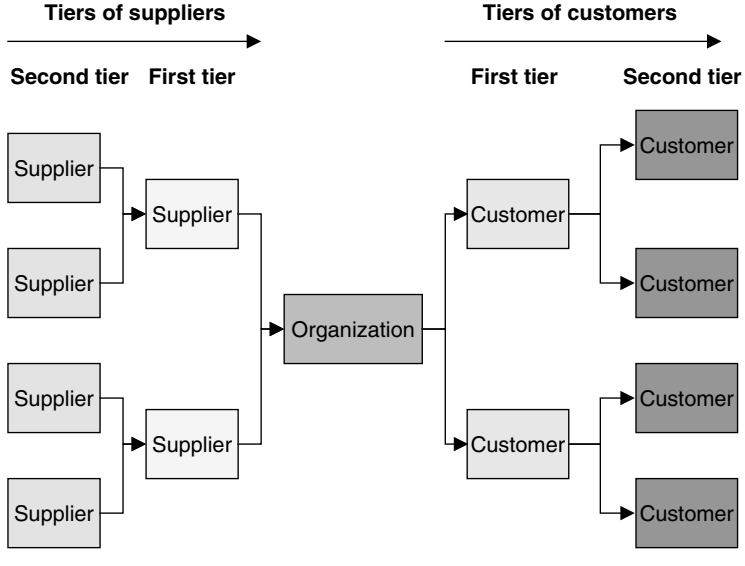


Figure 1.5 A simplified supply chain

observation suggests that the aggregate amount of stock held in a number of locations is:

$$AS(N_2) = AS(N_1) \times \sqrt{\frac{N_2}{N_1}}$$

where:

N_2 = number of planned future facilities
N_1 = number of existing facilities
$AS(N_i)$ = aggregate stock with N_i facilities

Worked example

AJT Transport of Manchester is planning to increase its services to mainland Europe. It currently has 12 depots with aggregate stock valued at £12 million and plans to expand to 16 depots. With a carrying cost is 20 per cent of value a year, what is the likely cost of this change?

Solution

We know that:

$$N_1 = 12 \text{ depots}$$

$$N_2 = 16 \text{ depots}$$

$$AS(N_1) = £12 \text{ million}$$

Then we can substitute these values to get:

$$AS(N_2) = AS(N_1) \times \sqrt{\frac{N_2}{N_1}} = 12 \times \sqrt{16/12} = £13.9 \text{ million}$$

The additional depots will raise stock holding costs by:

$$(13.9 - 12) \times 0.2 = 0.38 \text{ million or } £380{,}000 \text{ a year}$$

The best shape for a supply chain depends on many factors, such as the product's value, bulk, perishability, availability, profitability, and so on. It also depends on the organization's aims and business strategy. As a rule, a short, narrow supply chain gives an organization a lot of control over its logistics, but with a few, scattered intermediaries it is difficult to achieve either high customer service or low costs. Broadening the chain and adding more intermediaries gives higher customer service, but increases costs and reduces the organization's control. Making the supply chain long and narrow can reduce costs, but the organization loses some control and the customer service does not improve. Making the supply chain both

long and broad removes most control from the organization and raises costs, but gives good customer service. As you can see, organizations often have to find the best balance between costs and customer service. This is a common theme in inventory management, and we will return to it several times in the book.

Unfortunately, there is never a single 'best' shape for a supply chain, and managers have to look for designs that come closest to achieving their aims. One approach to doing this takes the following steps:

1. The logistics strategy sets the overall direction of logistics (as we shall see in the next chapter) so analyse this and find the aims of the supply chain.
2. Examine current operations, identify their failings and look for ways of overcoming these.
3. Design an outline structure for logistics, finding the number of facilities, best locations, modes of transport, investment in stocks, etc.
4. Make detailed plans, setting the size of each facility, stock holdings, material handling equipment, systems to develop, people to employ, transport needs, etc.
5. Get final approval from senior managers and agree the funding.
6. Finalize building designs, purchase land, choose contractors and build.
7. Finalize equipment design, choose equipment, suppliers and purchase.
8. Finalize systems design, for ordering, inventory control, billing, goods location, monitoring and all other systems.
9. Fit out facilities, install all equipment, systems, staff and test operations.
10. Open and receive stock, run final tests of all systems, finish training and begin operations.
11. Sort out teething problems and get things running smoothly.
12. Monitor and control, ensuring that everything works as planned, measure performance, revise targets, etc.

This is, of course, only a guideline to suggest the decisions in designing a supply chain. You can clearly see how decisions about the broader supply chain affect the stocks by, for example, setting the location, space available, handling facilities, systems and investment. You can also see how, conversely, attitudes towards stocks affect the design of the supply chain. If, for example, organizations accept that large stocks of finished goods must be kept near to customers, they will design supply chains to feed into these stocks. We will return to this theme in the next chapter.

Supply chain management is going through a period of rapid change. One clear trend is towards shorter chains, as organizations realize that they can both reduce costs and increase customer service by moving materials quickly through short chains. To achieve this, they remove layers of intermediaries and hold stocks in

fewer, larger facilities. Within the European Union, for example, efficient transport links mean that companies can replace a series of national warehouses by a single European logistics centre. Finding the best locations for these centralized stores can be very difficult. They might be near to factories, customers, transport, other facilities – or in areas with development grants. If an organization wants fast delivery, it has warehouses close to final customers; if it wants the lowest costs, it concentrates stocks in very large, centralized warehouses that are inevitably some distance from customers; if it imports and exports a lot of materials it might use warehouses near to ports, airports or rail terminals; if it manufactures goods, it has a stock of finished products near the factory. Here we cannot deal with details of the design of supply chain, but you should remember that these decisions include some of the key issues for inventory management. Again, we will discuss some of these in the next chapter.

Co-operation within a supply chain

A traditional view has each organization in a supply chain working largely in isolation and concerned only with its immediate suppliers and customers. This short-sighted view ignores the obvious point that the success – and survival – of the whole supply chain depend on its ability to satisfy final customers. Organizations in a supply chain increasingly recognize that they share a common overall objective, and should not compete with each other, but should co-operate to get final customer satisfaction. Competitors are not other organizations within the same supply chain, but are organizations in other supply chains. As Christopher (1996) says, 'supply chains compete, not companies'.

We can easily demonstrate the kind of problem that arises if organizations in a supply chain do not co-operate. Imagine a retailer who notices that demand for a product rises by 10 units in a week. When it is time to place the next order, the retailer assumes that demand is rising, and orders 20 extra units to make sure it has enough. The local wholesaler sees demand rise by 20 units, so it orders an extra 30 units to meet the growth. The regional wholesaler sees demand rise by 30 units, so it orders another 40 units. As this movement travels back through the supply chain, a relatively small change in final demand is amplified into a major variation for early suppliers. When demand from final customers moves down a bit, this is amplified into a collapse in demand for early suppliers.

Worked example

A simple supply chain has a manufacturer, regional and local wholesalers, a retailer and final customer. Each organization holds its own stock of one week's demand. In other words, each buys enough materials from its suppliers to make its closing stock at the end of the week equal to the demand during the week. Demand for a product is steady at 10 units a week. One week, however, demand from final customers rises to 20 units. Assuming that deliveries are very fast, how does this affect stocks in the supply chain?

Solution

The spreadsheet in Figure 1.6 shows these results. For each tier you can see:

- demand – which equals the amount bought by the following tier of customers;
- opening stock at the beginning of the week – which equals its closing stock in the previous week;
- closing stock at the end of the week – which must equal demand in the week;
- number of units bought – which equals demand plus any change in stock:

$$\text{buys} = \text{demand} + (\text{closing stock} - \text{opening stock})$$

In week 1 everything is going smoothly, with the usual 10 units flowing through the supply chain. Then in week 2 customer demand goes up to 20 units. The retailer must buy 20 units to meet this demand, plus an additional 10 units to raise its closing stock to 20 – so it buys 30 units from the local wholesaler. The local wholesaler has to supply this 30 units, plus an additional 20 units to raise its closing stock to 30 units – so it buys 50 units from the regional wholesaler.

	A	B	C	D	E	F	G	H	I	J
1	**Week**		**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
2										
3	**Customer**									
4		Demand	10	20	10	10	10	10	10	10
5										
6	**Retailer**									
7		Demand	10	20	10	10	10	10	10	10
8		Opening stock	10	10	20	10	10	10	10	10
9		Closing stock	10	20	10	10	10	10	10	10
10		Buys	10	30	0	10	10	10	10	10
11										
12	**Local wholesaler**									
13		Demand	10	30	0	10	10	10	10	10
14		Opening stock	10	10	30	30	20	10	10	10
15		Closing stock	10	30	30	20	10	10	10	10
16		Buys	10	50	0	0	0	10	10	10
17										
18	**Regional wholesaler**									
19		Demand	10	50	0	0	0	10	10	10
20		Opening stock	10	10	50	50	50	50	40	30
21		Closing stock	10	50	50	50	50	40	30	20
22		Buys	10	90	0	0	0	0	0	0
23										
24	**Manufacturer**									
25		Demand	10	90	0	0	0	0	0	0
26		Opening stock	10	10	90	90	90	90	90	90
27		Closing stock	10	90	90	90	90	90	90	90
28		Makes	10	170	0	0	0	0	0	0

Figure 1.6 Effect of varying demand in worked example

The regional wholesaler has to supply this 50 units, plus another 40 units to raise its closing stock to 50 units – so it buys 90 units from the manufacturer.

In week 3 we get the reverse effect as customer demand returns to 10 units. The retailer now reduces closing stock to 10 units, so it can meet all demand from stock and does not have to buy anything from the local wholesaler. This gives a demand of zero for all other tiers of suppliers. The stocks slowly returns to normal over following weeks, but a variation in customer demand of 10 units in one week, makes manufacturing jump by 160 units. The total amount of stock in the supply chain rises from 40 units to 190 units, and this will take 15 weeks to return to normal. Of course, we could criticize the inventory control policies here, but we will return to this theme in later chapters.

Any uncertainty in the supply chain – such as the jump in demand seen in the worked example – encourages organizations to hold higher stocks to give themselves a margin of safety. These extra stocks clearly increase costs. They also make the chain slow to react to changing conditions – when, for example, final customers start demanding a new product, all the stocks of old products in the supply chain have to be sold before the new ones appear. The way to avoid such problems is to co-ordinate the stocks and flow of materials. This brings a series of benefits, which include:

- lower costs – with lower stocks, less expediting, balanced operations, economies of scale, etc.;
- improved performance – with more stable operations, better planning, higher productivity of resources, etc.;
- improved material flow, with co-ordination giving faster and more reliable movements;
- better customer service, with shorter lead times and faster deliveries;
- more flexibility, with organizations reacting faster to changing conditions.

Christopher (1999) again summarizes the situation by saying that, 'Most opportunities for cost reduction and/or value enhancement lie at the interface between supply chain partners'.

Achieving co-operation in the supply chain

As you would imagine, there are many practical difficulties with achieving this co-operation. Many organizations simply do not trust other members of the supply chain, and they are reluctant to share information. Even with sufficient trust, there can be problems with different priorities, competition, data exchange, appropriate systems, skills, security, the complexity of systems, and so on. This raises the obvious question of how to achieve integration.

The first problem is overcoming the traditional view of organizations as adversaries. When an organization pays money to its suppliers, people assume that one

can only benefit at the expense of the other. If the organization gets a good deal, it automatically means that the supplier is losing out; if the supplier makes a good profit, it means that the organization pays too much. This adversarial attitude has major drawbacks. Suppliers set rigid conditions and, as they have no guarantee of repeat business, they see no point in co-operation and try to make as much profit from each sale as possible. At the same time, organizations have no loyalty, and they shop around to get the best deal and remind suppliers of the competition. Each is concerned only with their own objectives and will – when convenient to themselves – change specifications and conditions at short notice. The result is uncertainty about the number and size of orders, constantly changing suppliers and customers, changing products and conditions, different times between orders, no guarantee of repeat orders and changing costs.

To avoid these problems, organizations have to recognize that it is in their own long-term interest to replace conflict by co-operation. There are several ways of organizing this. The simplest appears when an organization has a good experience with a supplier and continues working with them over some period, developing a valuable working relationship. The key point with such informal arrangements is that there is no commitment. This has the advantage of being flexible and non-binding, but it has the disadvantage that either party can end the co-operation without warning and at any time that suits them.

Many organizations prefer a more formal arrangement, with a written contract setting out the obligations of each party. These are common when organizations see themselves as working together for some time, such as an electricity company agreeing to supply power at a fixed price for the next three years, provided a customer buys some minimum quantity. More formal agreements have the advantage of showing the details of the commitment, so that each side knows exactly what it has to do. On the other hand, they have the disadvantage of losing flexibility and imposing rigid conditions.

When an organization and a supplier are working well together, they may both feel that they are getting the best possible results and neither could benefit from trading with other partners. Then they might look for a long-term relationship that will guarantee their mutual benefits continue. This is the basis of a *strategic alliance* or *partnership*. Ellram and Krause (1994) prefer the term *supplier partnering* which they define as 'an ongoing relationship between firms, which involves a commitment over an extended time period, and a mutual sharing of information and the risks and rewards of the relationship'.

With such alliances, the supplier knows that it has repeat business for a long time, and can invest in improvements to products and operations; the organization knows that it has guaranteed – and continually improving – supplies. This often encourages suppliers to specialize in one type of product. They give such a commitment to the alliance that they reduce their product range, make these as efficiently as possible, and concentrate on giving a small number of customers a very high quality service. At the same time, customers reduce their number of suppliers, as they no longer need to look around to get the best deals. Japanese companies were among the first to develop strategic alliances, and at the time

when Toyota had formed partnerships with its 250 suppliers, General Motors was still working separately with 4,000.

Despite the clear evidence of benefits from co-operation in the supply chain, some people are not convinced, and say that each organization should independently pursue its own aims. You can imagine these different views in an example. Suppose a supplier has been delivering a product to an organization for some time, and has recently improved its operations to reduce the cost by 2 per cent. When it is time to negotiate this year's deliveries, what price should it quote? At one extreme is the view that the supplier has been working happily with the old level of profit, and should pass on all the savings in a lower price to make sure that it remains competitive and keeps the organization's custom. At the other extreme is the view that the supplier should maximize its own profits by keeping all the savings and increasing the price by (at least) the rate of inflation. In the middle is a compromise view which says the supplier should somehow share the benefits of lower costs with the organization. The final decision depends on aims, competition, power in the supply chain, and so on. Some formula for sharing the benefits is likely to give the best long-term results.

Summary

Inventory management can be viewed as one of the tasks of logistics. It is very closely related to other activities in the supply chain. As materials move through a supply chain, stocks are held at various points. The best results come when organizations within the same supply chain co-operate to ensure final customer satisfaction. There are several ways of achieving this co-operation, ranging from informal trading relationships through to partnerships.

Review questions

1.7 What is a supply chain?

1.8 Good customer service comes when stocks of finished goods are as close as possible to final customers. Do you agree with this?

1.9 Why should organizations within the same supply chain work together?

1.10 A company can only increase its profits by paying its suppliers less or charging its customers more. Do you think this is true?

Trends affecting stock

Organizations are constantly looking for ways of improving their operations and gaining a competitive advantage. We have seen how this encourages co-operation within a supply chain, and we mentioned some other issues, such as shorter supply chains, lower stocks, increasing customer service, and so on. In reality, these changes are taking place very quickly, and organizations are going through a period of rapid adjustment to the way they work. New practices and developments are making fundamental changes to operations, and many of these affect the role

and management of stock. We cannot look at all of these changes in detail, but will mention some of the most significant. The following list is, of course, nowhere near exhaustive and all kinds of changes are now affecting stocks.

Improving communications

Organizations are always introducing more sophisticated technology. Much of this affects the movement and storage of materials – with electronic identification of packages, satellite tracking of deliveries, automatic systems for moving goods, and so on. However, the greatest impact in recent years has come from improved communications. Consider, for example, the effect on purchasing materials. When a company wanted to buy something, it traditionally had to generate a description of the item, request for price, purchase order, order confirmation, contract terms, shipping papers, financial arrangements, delivery details, special conditions, invoices, and so on. All of these – and mountains of other paperwork – had to be printed and posted between organizations. By the 1990s technology had revolutionized these communications, with *electronic data interchange* or *EDI* allowing the direct exchange of data between remote computers. Supermarkets were among the first users of EDI, when they linked their stock control systems directly to suppliers' order processing systems. Then supermarket checkouts could record sales of each item, and when stocks got low the system automatically sent a message asking for another delivery. This use of *EPOS* – electronic point-of-sales – data gave less paperwork, lower transaction costs, faster communications, fewer errors, more integrated systems and closer business relations.

By 1997 about 2000 companies in the UK were using EDI (Stafford-Jones, 1997), and over the next few years electronic trading became more sophisticated and widespread. The mushrooming of e-mail was followed by all kinds of e-business, e-commerce – and soon 'e-anything'. The efficient transfer of information has been particularly useful for purchasing, which has developed into *e-purchasing* or *e-procurement*. This comes in many forms, with the two main versions based on *B2B* (business-to-business, where one business buys materials from another business) and *B2C* (business-to-customer, where a final customer buys from a business). By 2002 around 83 per cent of UK suppliers were using B2B (MRO Software, 2001), and the worldwide value of B2B trade was over US$2 trillion (The Gartner Group, 2002).

Two associated technologies supported EDI. The first is *item coding*, which gives every package of material moved an identifying tag. The tag is usually a bar code or magnetic stripe that can be read automatically as the package moves through the supply chain. Then stock control systems know where every package is at any time, and automatic materials handling can move, sort, consolidate, pack and deliver materials as needed. The second technology is *electronic fund transfer* or *EFT*, which automatically debits a customer's bank account and credits the supplier's. This completes a loop, with EDI to place orders, item coding to track the movement, and EFT to arrange payment.

Improved communications, together with better transport, mean that physical distances are becoming less significant. Organizations can become global in

outlook, buying, storing, manufacturing, moving and distributing materials in a single, world-wide market. As a result, international trade and competition continue to rise. Organizations used to look for competitors in the same town, but now they are just as likely to come from another continent. This trend is encouraged by free trade areas such as the European Union and the North American Free Trade Agreement.

Improving customer service

Customers have become increasingly knowledgeable about products and suppliers, and demand lower costs, higher quality and better service. In the past we might have gone to a local retailer to see what they were selling, but now we can surf the Web to compare the products offered by any company in the world.

To offset the increasing demands from customers, organizations are becoming more competitive. They might simply keep prices down, or they might find some other way of improving customer service. Both of these depend on the management of stocks. Low prices, for example, can only be achieved with low costs, and one significant factor here can be the cost of holding stock. It is normally in everyone's interests to make this as low as possible, so that customers pay as little as possible and the organization remains competitive. However, we cannot view the stockholding costs in isolation, as it is frequently linked to customer service. If an organization tries to reduce costs by holding less stock, it might find that there are more frequent shortages; or if it closes down a warehouse, the delivery time might rise. Although there is a continuing trend towards better customer service, we have to balance the gains this brings against the cost. As you will see in the rest of this book, finding the best balance between costs and service is a recurring theme in inventory management.

Concentration of ownership

Large companies can find economies of scale, and they have come to dominate many operations. There are, for example, many shops and transport companies, but the biggest ones continue to grow at the expense of small ones. The result is that most industries are dominated by a handful of major companies. This concentration is accompanied by changing power in the supply chain, with very large retail chains, such as Wal-Mart, Tesco, and Toys-Я-Us, demanding customized services from their suppliers.

At the same time there is concentration among suppliers, with fewer major companies. However, the effects are less clear here, as each customer is likely to work with a broader range of suppliers. The trend, though, is towards fewer suppliers with long-term commitments.

Outsourcing inventory management

Traditionally, each organization looked after its own logistics. Now, though, more organizations realize that they can benefit from using specialized companies

to take over part, or all, of their logistics, leaving them free to concentrate on their core activities. This is called *outsourcing* with the specialized companies being *third party operators*. A common starting point is to outsource transport to a specialized company. Next steps are to outsource warehousing, purchase of materials, materials handling, and other aspects of inventory management. This can bring significant benefits, which include the following:

- lower fixed costs, with organizations only paying for services they use;
- specialist suppliers who have expertise and use the best systems and practices;
- third parties can combine work from several customers to get economies of scale;
- guaranteed high, and agreed, levels of customer service;
- flexible capacity, dealing effectively with peaks and troughs in demand;
- lower exposure to risk from, say, varying demand;
- increased geographical cover and local knowledge;
- a convenient way of working in new markets.

Of course, there are disadvantages of outsourced inventory management, including reduced control, inability to respond to unusual circumstances, more complicated communications, conflicting objectives, less control over costs, and so on. Nonetheless, the advantages of third party operations are becoming clearer, with more organizations moving in this direction.

A variation on outsourcing has *vendor managed inventory* or *VMI* where suppliers manage both their own stocks and those held further down the supply chain. In this case the third party operator is an organization higher up the supply chain. This brings the usual benefits of outsourcing, together with much closer co-ordination and control of stocks. A slight variation on this has *co-managed inventory* where an organization and supplier somehow co-operate to manage stocks jointly.

Cross-docking

A traditional warehouse has materials delivered, it moves them into stock, keeps them until needed, and then delivers them to customers. *Cross-docking* co-ordinates the supply and delivery, so that materials arrive at the receiving area and are transferred straight away to the loading area where they are put onto delivery vehicles for customers. Ideally individual packages arrive and are passed on for delivery, but sometimes there is a limited amount of sorting, perhaps breaking down larger deliveries into smaller amounts for each customer. Some people prefer the term *flow through stock* when there is this kind of sorting.

A related arrangement uses *drop-shipping*, where materials do not actually go to the warehouse, but are delivered directly from upstream suppliers to downstream customers. If, for example, you buy a Hotpoint washing machine from a Dixon's store, you will probably find that the machine is delivered directly

from Hotpoint without ever visiting a Dixon's warehouse. Methods like this are becoming increasingly common, as more customers buy through the Web, or find other ways of trading with earlier parts of the supply chain, such as mail order, telephone shopping or buying directly from manufacturers. This has the benefits of reducing lead times, reducing costs to customers, having manufacturers talking directly to their final customers, allowing customers access to a wider range of products, and so on.

Any method that avoids putting materials into stock at a warehouse can dramatically reduce stock levels and associated administration. In the extreme, the only stock is within delivery vehicles, giving *stock on wheels*. This is not, however, a solution for all problems, as the co-ordination can be difficult, and this needs a certain volume of trade to stop the small, frequent deliveries becoming prohibitively expensive. Many parcel delivery services such as FedEx, UPS, Omega and DHL have grown as a way of giving rapid delivery of small amounts of materials at reasonable cost.

Postponement

Manufacturers typically move finished goods out of production and keep them in a store of finished goods until needed. When there are many variations on a basic product, this can give high stocks of similar products. Postponement moves almost finished products into stock, and delays final modifications or customization until the last possible moment. You can imagine this with 'package-to-order', where a company keeps a product in stock, but only puts it in a box written in the appropriate language when it is about to ship an order.

Manufacturers of electrical equipment, such as Phillips and Hewlett-Packard, used to build into their products the transformers and plugs needed for different markets. Then they had to keep separate stocks of products destined for each country. Now they make the transformer and cables as separate, external units. They only keep stocks of the basic, standard products, and customize them for different markets by adding the proper transformers and plugs at the last minute. The result, of course, is much lower stocks.

Increasing environmental concerns

One cultural trend that is affecting inventory management is the growing concern about air pollution, water pollution, energy consumption, urban development and waste disposal. It is fair to say that the whole area of logistics does not have a very good reputation for environmental protection – demonstrated by the emissions from heavy lorries, use of greenfield sites for warehouses, calls for new road building, use of extensive packaging, ships illegally flushing their fuel tanks, oil spillages from tanker accidents, and so on. On the positive side, though, logistics is moving towards 'greener' practices, with more energy-efficient vehicles, control of exhaust emissions, reuse of packaging, switching to environmentally friendly modes of transport, increasing recycling through reverse logistics, added safety features to ships, development on brown-field sites, and so on. There is increasing

recognition that careful management can bring both environmental protection and lower costs.

Summary

Organizations are continually looking for ways to improve their operations. The results have created a series of trends that affect inventory management. Among these are increasing use of technology, improved customer service, global operations, outsourcing, new arrangements for stock management, and so on.

Review questions

1.11 Why do you think that there are currently such rapid changes affecting inventory management?

1.12 Are the effects of these trends independent of each other?

1.13 Several trends in operations are leading to lower stocks. Is it inevitable that stocks will continue to decrease?

Changes to aggregate stocks

Changing views of stock

The trends outlined in the last section show that organizations are changing their attitudes towards stock. This is not new, but is part of a continuing pattern. For most of history, stocks have been considered measures of wealth or well-being and were, therefore, beneficial. The family with the biggest store of food was least likely to starve, the company with most raw materials was insulated from shortages, and the business with most money in the bank was most secure. The obvious conclusion was that stocks should be maximized, as this gives the greatest benefit.

At times when the production and distribution of any material are uncertain, it certainly makes sense to avoid problems by collecting as much stock as possible. However, by the turn of the twentieth century industrialized countries had more or less secure supplies of most materials. The uncertainty in supply was greatly reduced, and this brought a new attitude towards stock. Organizations could now buy materials when they were needed – rather than when they were available – and they looked for more rational ways of controlling stock levels. This new outlook suggested that stocks were expensive and needed formal management. In particular, organizations should look for ways of minimizing some aspect of cost. Sometimes people were sidetracked from this aim and in the 1920s, for example, there was a craze for minimizing stocks rather than costs. Unfortunately, many companies hit problems when they reduced stocks to levels that made it impossible to work effectively or maintain any kind of customer service.

In the late 1920s *scientific inventory control* became the main approach to inventory management, using mathematical models to find optimal stock levels. As

you can imagine, a continuing problem has been to find agreement about what 'optimal' levels might really be. Are these the levels that minimize total cost, or some aspect of cost, or give high customer service, or minimize stock-outs and disruptions, or give highest return on investment, or maximize stock turnover, or achieve one of dozens of other measures of performance?

For some time, it was felt that a 'fixed accelerator' could define an optimal stock level as some fixed proportion of sales (for example, Abramovitz, 1950). In practice, this proved ineffective and there was a move towards a 'flexible accelerator' to allow for differences between aims and actualities, time delays, and so on (for example, Lovell, 1961, 1964). Unfortunately, this approach also had its failings, and a more flexible approach was developed, which used a range of models to deal with different circumstances. These models grew increasingly sophisticated, and remained the main approach to inventory management for most of the last century. We will describe some of these models in Chapters 3 to 5.

More recently, new ideas have emerged about inventory management. These do not look for the best policy for dealing with uncertainty in supply and demand, but look for ways of removing the uncertainty. When there is no uncertainty, the stocks can be eliminated, or at least minimized. We will describe approaches of this kind in Chapters 9 and 10.

To summarize this brief review, for most of history, stocks were seen as beneficial and organizations attempted to maximize their holdings; in the last century, organizations realized that stocks were expensive and looked for policies that defined optimal stock levels; most recently organizations have looked for ways of eliminating stocks. This is, of course, only a broad overview and you can find many organizations that work with very high stocks. Shops, for example, keep high stocks so that customers can see a range of goods and do not have to wait for deliveries. In the same way, there are many parts of the world where supplies are still not reliable and when materials become available, organizations buy as much as possible.

Aggregate national stocks

In recent years, organizations have been working to lower stocks without affecting either their own efficiency or customer service. Perhaps we should look for some broad evidence to see how successful they have been. Surveys give some evidence for success, with the Institute of Grocery Distribution finding that stock levels in retail distribution centres fell by 8.5 per cent in the year to 1998 (Institute of Grocery Distribution, 1998), and the Institute of Logistics finding that some UK companies had 'managed to almost halve the stockholding requirements since the 1995 survey' (Institute of Logistics, 1998).

More general evidence comes from government statistics. In the UK aggregate stock holdings are about £100,000 million, divided roughly equally between raw materials, work in progress and finished goods (Office of National Statistics, 2002). There are surprisingly wide changes in this aggregate national stock. Some changes are planned, as organizations adjust their stocks, while others are a consequence of broader economic influences, when, for example, the economy

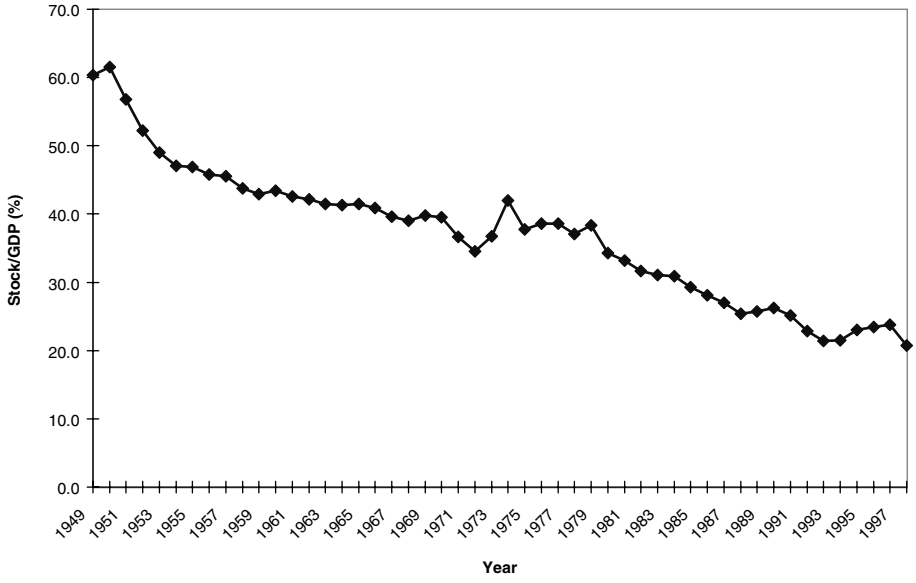


Figure 1.7 Aggregate stock as a percentage of GDP for the UK

declines, sales fall and organizations are left with higher than expected stocks of unsold products. If we consider the aggregate national stock as a proportion of gross domestic product, we get a useful measure that overcomes some effects of changing economic conditions and focuses on the changes which are positively planned. Figure 1.7 shows this result for the second half of the last century (Central Statistics Office, 1966–1983, and 1984–1996).

The figures show a clear pattern. At the end of the 1940s and into the early 1950s there was a rapid decline in stocks as the economy returned to normal after the Second World War. From the early 1950s to the early 1970s there was a steady decline, which gives evidence for improving inventory management. In the early 1970s there was some disturbance caused by a rapid increase in the price of oil, and the economic disruption that followed. At this time, the costs of raw materials rose sharply and there were frequent shortages, while declining sales left finished goods unsold. After this disturbance, the long-term trend continued, with organizations improving their operations and working with ever-lower stocks.

Effects of the business cycle

Aggregate stock holdings are clearly influenced by general economic conditions over which individual organizations have no control. We can illustrate one aspect of this by business cycles. A traditional view of business cycles starts with industry being optimistic about the future. Sales are expected to rise, so production increases to match perceived future demand and the economy expands. Actual sales lag behind this increased production, so there is a build-up of

stock. At some point, industry loses confidence and cuts back on production to reduce the excessive stocks, and the economy contracts. This recession – or at best stagnation – continues until stocks are lower, production is not meeting expected demand, and industry again expands (as shown in Figure 1.8).

Nobody has found a precise cause or explanation for business cycles, and there is a general belief that each cycle is in some ways unique. It is, however, widely accepted that long-term business cycles and stocks are closely related, and that stock levels – as one of the easiest factors to change – tend to fluctuate more than the business cycle itself. One view has variations in stocks as actually causing business cycles. Klein and Popkin (1961) suggest that controlling 75 per cent of the variation in stock levels in the United States between the World Wars would have avoided all recessions. Such findings might encourage governments to prohibit wide fluctuations in stocks by taxes or other means. In practice, such measures have never been tried, mainly because of the difficulties in defining 'excessive' fluctuations and finding a reasonable way of preventing them.

As well as the general economic climate, there are specific reasons for aggregate stock to vary. It is, for example, sometimes suggested that high interest rates should lead to low stocks, as it becomes more expensive to finance them. This argument is not really convincing, as finance is only one of the costs of inventory

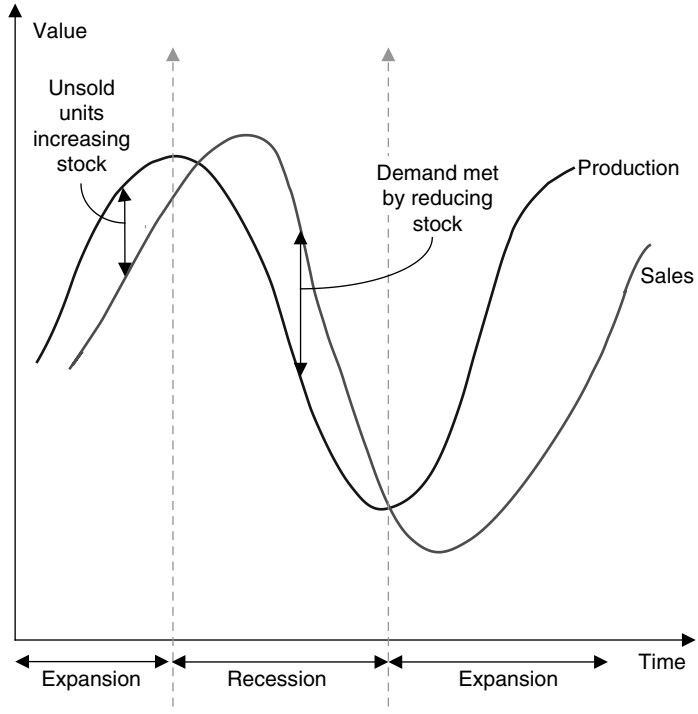


Figure 1.8 Stock levels in a business cycle

and the other costs are so high that a small variation in interest rates should have little effect. More significantly, an organization that can work with lower stocks when interest rates are high should always work with lower stocks and reduce its on-going costs.

A more convincing argument shows the effect of inflation. At times of high inflation stock levels rise as organizations buy more materials at the current lower prices to avoid the higher prices that are likely with future orders. At the same time, the book value of stocks increases, raising the value of assets and making stocks more attractive.

Summary

Attitudes to stocks have changed over time. The current trend is towards lower stocks. This can be seen at a national level, where aggregate stocks show a long-term decline as a proportion of Gross National Product. National stock levels are also affected by business cycles, but the details of this relationship are unclear.

Review questions

1.14 Why are organizations moving towards lower stocks?

1.15 Is the objective of minimizing the costs of stock holding the same as minimizing stock?

1.16 Why do stock levels tend to fall during periods of recession?

1.17 On a national scale, reducing variations in stock levels would reduce the severity of business cycles. Do you think this is true?

Chapter review

- This chapter introduced the ideas behind inventory management, laying the foundations for later chapters. It started by defining some important terms.
- Every organization holds stock of some kind. There are many different materials held and arrangements for storage, but they all need careful inventory management.
- The main purpose of stocks is to act as a buffer between operations. They allow operations to continue normally through variations and uncertainty in supply and demand.
- Stocks are held at various points in their supply chains. In the past these stocks have largely been considered as independent, but there are clear advantages in co-ordinated management.
- Organizations are going through a period of considerable change. Many of these changes have direct effects on stocks. These include trends towards higher technology, improved customer service, global operations, and so on.

- Partly as a result of changing operations, organizations have changed their view of stocks. The current view is that they are expensive and should be reduced to the minimum level that can give acceptable customer service.
- This view has had an effect on aggregate, national stocks, which have been in a long-term decline.

Project

The purpose of this project is to get a view of stock holdings at a national level. Figure 1.7 showed the ratio of aggregate stocks to GDP for the UK during the second half of the twentieth century. Do these figures really show that organizations are deciding to work with lower stock levels or are there other explanations?

Update the figures and see if the trend is continuing. Are there any other factors that could be influencing recent figures? Find comparable figures for other countries. How do these compare with the UK? How can you explain any significant differences?

Discussion questions

1.1 Is it true that every organization holds stock?

1.2 Organizations hold stock to give a cushion between operations. But this stock is expensive, so a better approach would solve any problems and do away with the need for this buffering. Does this seem a reasonable suggestion? If it is, how could we do away with the need for buffers?

1.3 Organizations in a supply chain can never really co-operate, as they compete for available money. Customers should use every available means to pay the lowest price for materials; suppliers should charge the highest prices they can. Is this a more realistic view of relationships in a supply chain?

1.4 Many trends in business have a direct impact on stock management. What do you think are the most important trends at the moment?

1.5 Stock levels are inevitably declining. Eventually we will be able to work without any stock at all. Do you think this is true?

1.6 Why have the stocks in some countries fallen faster than in other countries?

References and further reading

Abramovitz, M. (1950) *Inventories and Business Cycles*. New York: The National Bureau of Economic Research.

Central Statistics Office (1966–1983) *National Income and Expenditure*. London: HMSO.

Central Statistics Office (1984–1996) *United Kingdom National Accounts*. London: HMSO.

Christopher, M. (1996) Emerging issues in supply chain management, *Proceedings of the Logistics Academic Network Inaugural Workshop*, Warwick.
Christopher, M. (1999) Global logistics: the role of agility, *Logistics and Transport Focus* 1(1).
Ellram, L.M. and Krause, D.R. (1994) Supplier partnerships in manufacturing versus non-manufacturing firms, *The International Journal of Logistics Management*, 5(1), 43–53.
The Gartner Group (2002) website at *www.gartner.com.*
Institute of Grocery Distribution (1998) *Retail Distribution 1998*. Herts: IGD.
Institute of Logistics (1998) *European Logistics: Comparative Survey*. Corby: Institute of Logistics.
Klein, L.R. and Popkin, J. (1961) An economic analysis of the post war relationship between inventory fluctuation and change in aggregate economic activity, in *Inventory Fluctuation and Economic Stabilization*. Washington, DC: The Joint Economic Committee.
Lovell, M.C. (1961) Manufacturers' inventories, sales expectations and the accelerator principle, *Econometrica*, 29, 293–314.
Lovell, M.C. (1964) Determinants of inventory investment, in I. Friend (ed.). *Models of Income Distribution*. Princeton, NJ: Princeton University Press.
MRO Software (2001) *Supplying the Goods*. London: MRO Software.
Office of National Statistics (2002) *UK Economic Accounts*. London: HMSO.
Stafford-Jones, A. (1997) Electronic commerce: the future with EDI, *Logistics Focus*, 5(9), 9–10.
Waters, D. (ed.) (2003) *Global Logistics and Distribution Planning*, London: Kogan Page.

Websites

Many websites describe some aspects of inventory management, with the following giving useful starting points

www.apics.org American Production and Inventory Control Society
www.inventorymanagement.com Centre for Inventory
www.cris.com Inventory Control Forum
www.poms.org Production and Operations Management Society
www.iomnet.org.uk Institute of Operations Management
www.theorsociety.com Operational Research Society

2

Stocks within an Organization

Aims of the chapter

The last chapter gave a general introduction to inventory management. In this chapter we are going to look more specifically at the role of stock within an organization. We will see how stock allows an organization to work efficiently, achieve its broader aims, and maintain an acceptable level of customer service. The chapter describes the context for inventory management within an organization and introduces some of the methods.

After reading this chapter you should be able to do the following:

- show how inventory management fits into the broader function of logistics;
- explain how the aims of inventory management are set within an organization;
- appreciate the strategic importance of stock;
- describe the costs associated with holding stock;
- list some basic questions of inventory control;
- outline some different approaches to stock control.

This chapter emphasizes:

- *logistics*, which moves materials through the supply chain;
- *stock* or *inventory management*, which includes all the decisions made about stock within an organization;
- *aims* and *approaches* of inventory management.

Inventory management and logistics

Functions of logistics

We saw in the last chapter that logistics is responsible for all aspects of the movement and storage of materials in a supply chain, while inventory management is specifically responsible for the storage of materials. Rather than viewing inventory

management as a separate function we could, therefore, describe it as one of the activities that combine to form the broader function of logistics. All of these activities are so closely related that it would certainly be difficult to draw a boundary between inventory management and, say, purchasing, materials handling or warehousing.

But if inventory management is so entwined with the other activities of logistics, we should clearly say what these other activities are. Depending on circumstances, an organization might include many different activities within its logistics, and the most common are given in the following list.

- *Procurement or purchasing* which usually initiates the flow of materials through an organization by sending a purchase order to a supplier. This means that procurement has to find suitable suppliers, negotiate terms, set conditions, organize delivery, arrange insurance, authorize payment, and do everything needed to get materials into the organization. In the past, this has been a largely clerical job centred on order processing, but now it is recognized as being responsible for major expenditure and giving a critical link with suppliers. The whole area is undergoing major changes as a result of e-commerce.
- *Inward transport or traffic* which actually moves materials from suppliers to the organization's receiving area. Important decisions concern the mode of transport (road, rail, air, etc.), policies for outsourcing, choice of transport operator, route, safety and legal requirements, timing of deliveries, costs, etc.
- *Receiving* makes sure that materials delivered correspond to the order, acknowledges receipt, unloads delivery vehicles, inspects materials for damage, and sorts them.
- *Material handling* moves materials from receiving and puts them into stores. It is responsible for all movements of materials within an organization, so it also removes materials from stores, takes them to the place they are needed, and generally moves materials between operations. Its aim is to provide efficient movements, with short journeys, using appropriate equipment, with little damage, and using special packaging and handling where needed.
- *Warehousing or stores* takes care of materials held in stock until they are needed. Warehousing makes sure that materials have the right conditions, treatment and packaging to keep them in good condition, and are available quickly when needed.
- *Stock or inventory control* sets the overall policies for stock, considering the materials to store, investment, customer service, stock levels, order sizes, order timing and so on.
- *Order picking* finds and removes materials from stores. Typically materials for a customer order are located, identified, checked, removed from racks, consolidated into a single load, wrapped and moved to a departure area for loading onto delivery vehicles.

- *Outward transport* takes materials from the departure area and delivers them to customers (with concerns that are similar to inward transport).
- *Physical distribution* is a general term for the activities that deliver finished goods to customers. It is often aligned with marketing and forms an important link with downstream activities.
- *Recycling, returns and waste disposal* There are sometimes problems with delivered materials – perhaps faults, the wrong materials or too much delivered – and logistics has to collect them and bring them back to the supplier. This *reverse logistics* also collects associated materials such as pallets, delivery boxes, cable reels and containers (the standard 20-foot-long metal boxes that are used to move goods). Some materials are brought back for recycling (such as metals, glass, paper, plastics and oils) while others are brought back for safe disposal (such as dangerous chemicals).
- *Location* finds the best site for those activities that can be in different places. Stocks of finished goods, for example, can be held at the end of production, moved to nearby warehouses, sent to large centralized facilities, put into stores near to customers, passed on to be managed by other organizations, or a range of alternatives.
- *Communication* Alongside the physical flow of materials is the associated flow of information. This links all parts of the supply chain, passing information about products, customer demand, materials moved, timing, stock levels, availability, problems, costs, service levels, and so on. Co-ordinating the flow of information can be very difficult, and logistics managers often describe themselves as processing information rather than moving goods. Christopher (1996) supports this view by saying that, 'Supply chain competitiveness is based upon the value-added exchange of information'.

These activities are all closely related and often overlap. We should not, therefore, try to draw artificial boundaries between them, but should view them as different aspects of a single logistics function. This reinforces the point that all logistics activities have a direct impact on the stocks. To take one example, we can look at the mode of transport, which describes the type of delivery vehicle used. The alternative modes are road, rail, air, inland waterway, ocean shipping or pipeline – and the choice is determined by locations, infrastructure, weight and volume carried, value of goods, customer service offered, urgency and a series of other factors. But this choice has a direct effect on stocks. Air travel is fastest and most expensive while shipping is slowest and cheapest – so the choice affects lead times and delivery costs. It also affects the amount of pipeline stocks (which are the materials in transit between locations), with airfreight having little stock on short journeys and shipping having weeks of stock out at sea. We could go on with more examples of the links between inventory management and other activities of logistics, but the main point is obvious. Inventory management is part of a broader function, and when we focus on this one activity, we should not separate it from the wider context.

Integrating logistics within an organization

Figure 2.1 shows how the activities of logistics fit into an organization. Traditionally, each activity has been managed separately, so there might be a distinct Purchasing Department, Transport Department, Warehouse, Distribution Fleet, and so on. Unfortunately, this fragmentation brings a number of problems. Suppose, for example, that warehousing tries to save money by reducing the stock of raw materials. This might give more frequent shortages, disrupt operations, increase expediting, raise the costs of emergency orders and special deliveries, and generally reduce the level of service. The cost of warehousing is lower, but overall cost to the organization is considerably higher. In the same way, an independent purchasing department might reduce its administrative costs by sending fewer, larger orders to suppliers but this increases stock levels and raises the amount of money tied up in the warehouse.

This kind of conflict is almost inevitable when logistics is divided into separate activities. Each part has different objectives, and there is duplicated effort and wasted resources. To put it briefly, fragmented logistics has the disadvantages of:

- creating different, often conflicting, objectives within an organization;
- duplicating effort and reducing productivity;
- interrupting information flows;
- reducing co-ordination;
- increasing uncertainty;
- making planning more difficult;
- introducing unnecessary buffers, such as stocks of work-in-progress;

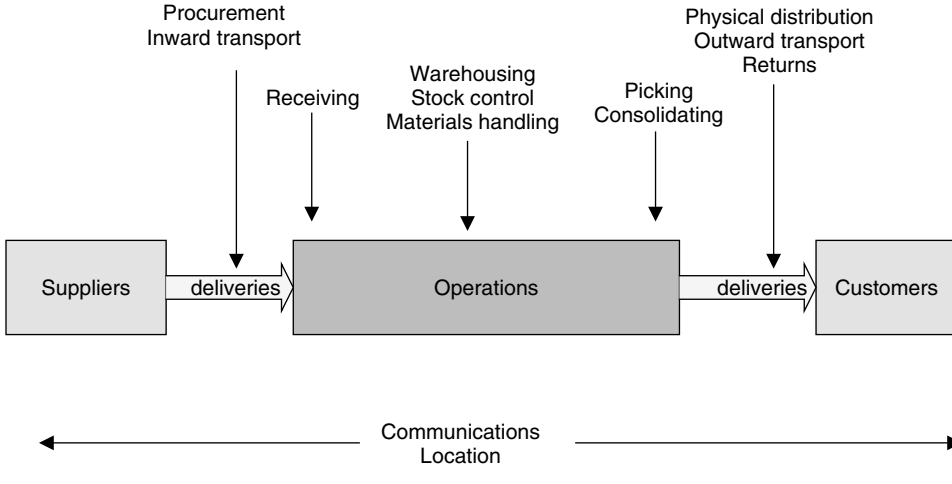


Figure 2.1 Place of logistics activities within an organization

- obscuring important information, such as the total costs involved;
- giving logistics a low status.

We can avoid these problems by considering logistics as a single, integrated function. Decisions should then be made to get the best results for the whole function, rather than for any single activity. When we make decisions about inventory management we should look at their broader impact and aim for the best overall result.

Worked example

RP Turner Corporation makes pipeline valves for the oil industry in Western Canada. It buys materials from Japan, the USA and Eastern Canada, manufactures valves in Edmonton, Alberta, and ships the finished products to oil fields in the North. In 2000 it started a major project to reduce the cost of logistics.

Managers soon found that the separate activities worked more or less independently. This was sometimes all too obvious when the three main Departments – Marketing, Production and Finance – were in different locations. Production was in Edmonton, as the nearest major city to the oil fields; Marketing was in Calgary near to oil company headquarters; Finance (including procurement) was in Vancouver near the port and financial centre. Canada is a big country, so Production was a thousand kilometres away from Finance, 500 kilometres away from Marketing and over two thousand kilometres from delivery points. How would you set about improving logistics?

Solution

One obvious step is to integrate the different activities of logistics. The company started this and found that each activity had different – and often conflicting – aims.

Marketing wanted:

- high stocks of finished goods to satisfy customer demands quickly;
- a wide range of finished goods always held in stock;
- locations near to customers to allow delivery with short lead times;
- production to vary output in response to customer orders;
- emphasis on an efficient distribution system;
- optimistic sales forecast to ensure production has enough capacity for actual demands.

Production wanted:

- high stocks of raw materials and work in progress to safeguard operations;
- a narrow range of finished goods to give long production runs;
- locations near to suppliers so that they can get raw materials quickly;
- stable production to give efficient operations;
- emphasis on the efficient movement of materials through operations;
- realistic sales forecasts that allow efficient planning.

Finance wanted:

- low stocks everywhere;
- a narrow range of products to give low unit costs;
- few locations to give economies of scale and minimize overall costs;
- long production runs to reduce unit costs;
- make-to-order operations;
- pessimistic sales forecasts that discourage under-used facilities.

Over the next two years the company integrated and centralized its operations at the main plant in Edmonton. This both reduced costs and improved customer service.

Summary

Inventory management can be viewed as a part of the broader logistics function. Decisions about stocks are linked, in some way, to all other areas of logistics. It is best to consider a single, integrated logistics function, and look for the inventory policies that give the best overall results.

Review questions

2.1 What activities are normally included in logistics?

2.2 Every activity should concentrate on its own objectives, and not interfere with others. This will always give the best results. Do you agree with this?

Setting the aims of inventory management

Levels of decision

We have already mentioned – at least implicitly – some of the aims of inventory management. We know, for example, that managers want to guarantee that

materials are available when they are needed, and that costs are controlled. In a broader sense we have really suggested three types of objectives for inventory management. The first takes a very broad view, and has inventory management contributing to the smooth flow of materials through the entire supply chain. The second takes an organizational view and has inventory management supporting logistics in achieving the overall aims of the organization. The third takes a functional view and has inventory managers making sure that materials are available when they are needed.

Most managers work with this kind of hierarchy of goals and decisions. If we look at the decisions made in any function, some are very important with consequences felt over many years, while others are less important, with consequences felt over days or even hours. We can use their importance to classify decisions as follows.

- *Strategic decisions* are most important, have effects over the long term, use many resources and are the most risky. These set the overall direction for operations.
- *Tactical decisions* are concerned with implementing the strategies over the medium term; they look at more detail, involve fewer resources and some risk.
- *Operational decisions* are concerned with implementing the tactics over the short term; they are the most detailed, involve few resources and little risk.

A traditional view has senior managers making the strategic decisions that set their organization on its course. These strategic decisions give the objectives, constraints and context for the tactical decisions made by middle managers. These, in turn, give the objectives, constraints and context for operational decisions made by junior managers. New styles of management and improved technology have changed this traditional pattern, and now we rarely see such a strict hierarchy. Most decisions are discussed, negotiated and agreed rather than simply passed down, and there is a growing recognition that the best person to make a decision is the person most closely involved with it.

There are several types of strategic decision (illustrated in Figure 2.2). People use different names for these, but they commonly start with a *mission* or *vision*, which gives a statement of the organization's overall aims. Then a *corporate strategy* shows how a diversified corporation will achieve its mission, and a *business strategy* shows how each business within the diversified corporation will contribute to the corporate strategy. Below this are *functional strategies* which describe the strategic direction of each function, including logistics. These functional strategies lead to tactical and operational decisions within each function.

This structure shows how the aims of inventory management are set. The higher strategies set the goals and overall direction; the logistics strategy shows how logistics will support these higher aims. Then the tactical and operational decisions show how the logistics strategy is implemented. Decisions for inventory management come at all levels, forming part of the logistics strategy and the lower tactical and operational decisions. Suppose, for example, an organization has a business strategy of being the lowest cost provider of some product; the logistics

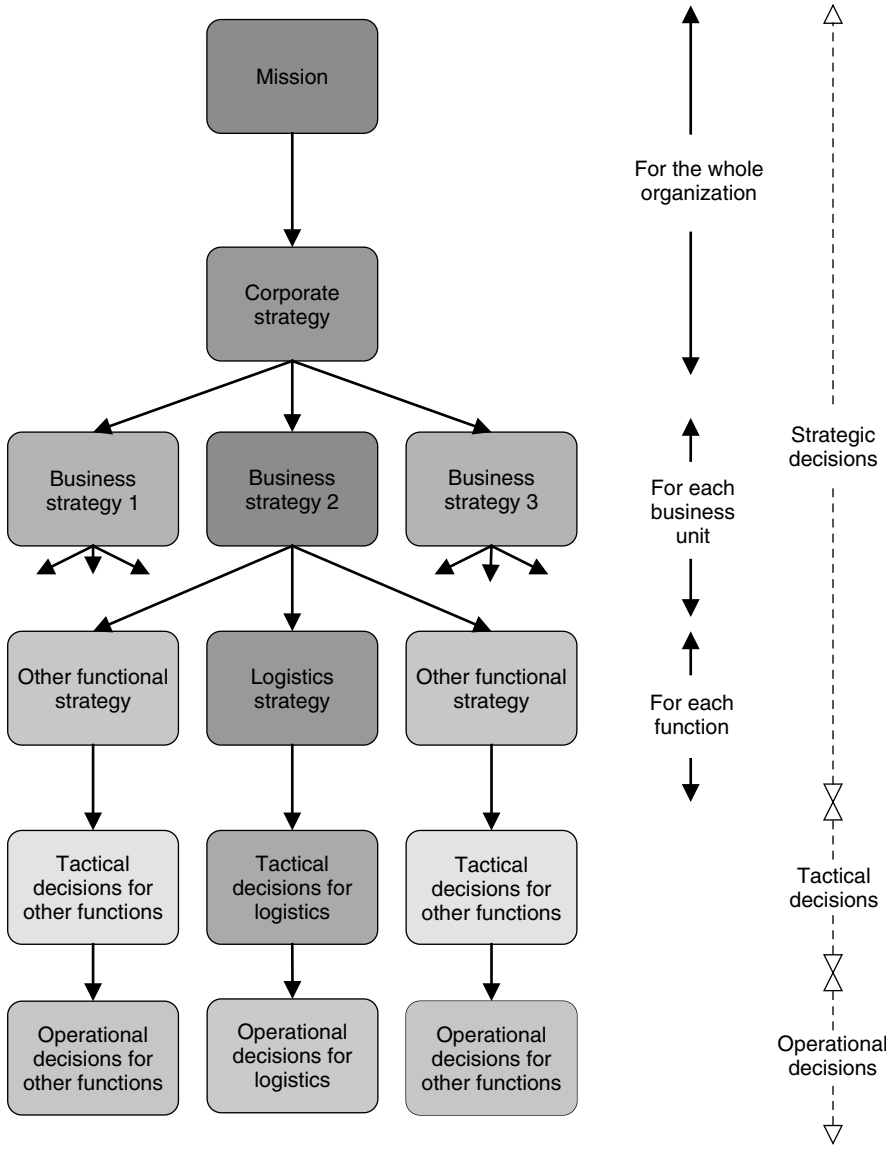


Figure 2.2 Types of decision within an organization

strategy will support this by reducing logistics costs to a minimum, and inventory managers will look for the lowest possible costs of holding stock. If the business strategy includes fast deliveries to customers, the logistics strategy defines policies of having stocks located near to customers, and inventory management makes the decisions to implement this. A strategic decision to sell products through a website leads to tactical decisions about warehousing, investment in stock, transport, materials handling, recruiting and training, customer service, and so

on; these tactical decisions, in turn, lead to operational decisions about resource scheduling, inventory control, expediting, vehicle routes, etc.

This view of decisions is obviously simplified and there is rarely such a rigid 'top down' approach. The distinctions between strategic, tactical and operational decisions are not this clear. Stock levels, for example, are a strategic issue for decisions about building a warehouse for finished goods or shipping directly to customers, a tactical issue when deciding how much to invest in stock, and an operational issue when deciding how much to order this week. Customer service is a strategic issue when designing the supply chain, a tactical issue when organizing transport for delivery, and an operational issue when scheduling the next delivery. The important point is that the strategy leads to a whole series of related decisions at different levels, and decisions about stocks have to fit into – and support – this overall picture (shown in Figure 2.3).

Alternative strategies

The business strategy sets the scene for all decisions about stocks, but what are the alternative strategies? Essentially a business strategy shows how the organization

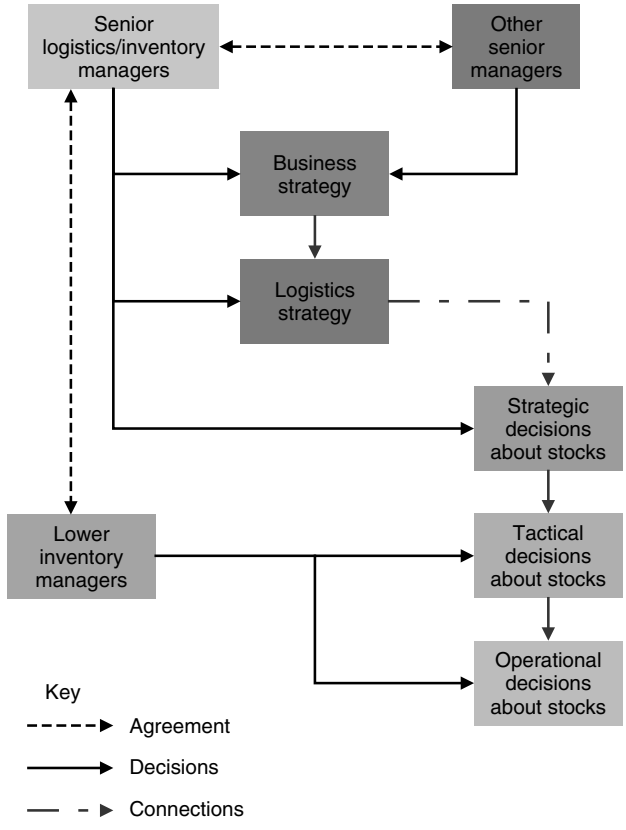


Figure 2.3 Structure of decision about stocks

can satisfy customers – for without satisfied customers no organization will survive in the long term, let alone make a profit, increase shareholder value, or achieve any other measure of success. Michael Porter (1985) suggested three basic strategies:

- *cost leadership* makes the same, or comparable, products more cheaply;
- *product differentiation* makes products that customers cannot get anywhere else;
- *niche suppliers* find a unique niche in the market.

The two dominant business strategies of cost leadership and product differentiation lead to *lean* and *agile* strategies for the supply chain. A lean strategy tries to do every operation with the least possible resource – people, space, stock, equipment, time, money, etc. This seems an obvious approach, but in 1970 Robert Townsend said that, 'All organizations are at least 50 per cent waste – waste people, waste effort, waste space and waste time'. A lean strategy looks for ways of eliminating this waste. The typical approach does a detailed analysis of current operations, and then removes operations that add no value, eliminates delays, simplifies movements, reduces complexity, increases efficiency, finds economies of scale, and so on.

Toyota carried out much of the early work on lean operations (see Ohno, 1988, and Womack *et al*., 1990) and they identify six areas where waste is most likely to occur (Monden, 1983):

- *Quality* – that is too poor to satisfy customers.
- *Wrong production level or capacity* – making products that are not currently needed or having unused capacity.
- *Poor process* – having unnecessary, too complicated or time-consuming operations.
- *Waiting* – for operations to start or finish, for materials to arrive, for equipment to be repaired, etc.
- *Movement* – with products making unnecessary, long, or inconvenient movements during operations.
- *Stock* – holding too much stock, increasing complexity and raising costs.

High stock levels are often seen as wasteful in themselves, but are also symptoms of other more pervasive waste within the organization. Over-production will obviously give high stocks of finished goods, while poor procurement can give high stocks of raw materials. The effects are, however, usually more subtle than this. Bottlenecks in a process, or processes that are too long, give higher stocks of work in progress; unreliable suppliers encourage higher stocks of raw materials; poor delivery schedules to customers give high stock of finished goods; poor quality materials encourage higher stocks throughout the organization. The lean approach says that reducing stock saves money directly, and it also highlights

problems that are currently hidden. Once these problems are identified, there is an incentive to solve them and improve overall performance.

Some people say that a lean strategy puts too much emphasis on costs, and cannot deal with changing conditions, increasing competition, or more sophisticated and demanding customers. The alternative is an agile strategy that gives a high customer service by responding quickly to different or changing circumstances. In practice, there are two aspects of agility. First, there is the speed of reaction; agile organizations keep a close check on customer demands and react quickly to changes. Second, is the ability to customize or tailor products to individual customer demands. These are, of course, different aspects of customer service, and the implication is that customer satisfaction is a prime concern, even if this comes at somewhat higher price.

At first sight, the aims of lean and agile operations seem contradictory. One looks to minimize costs, and sees customer service as a constraint; the other looks to maximize customer service, and sees costs as a constraint. In practice, of course, there is not such a clear divide between the two strategies. If, for example, a supplier improves EDI links with its customers, it can both reduce costs and increase customer service – becoming both leaner and more agile. Both strategies recognize that customer satisfaction and low costs are dominant themes, and there must be some balance between the two. The main difference is the point where this balance is found.

Strategic focus

Under the umbrella of lean and agile strategies, organizations can focus on specific features of their products and operations. For example, an organization with an overall agile strategy might focus on its ability to deliver goods very quickly to customers. There are many possible areas for focus, including the following:

- *Timing* – fast deliveries give good customer service, and they also bring lower costs (by having less stock in the supply chain, less expediting, etc.), improved cash flow (by not having to wait so long for payment), less risk (by reducing changes to orders, obsolete stock, etc.) and simpler operations (by eliminating delays and unnecessary stores).
- *Quality* – guaranteeing high quality goods and services.
- *Product flexibility* – which is the ability to customize products to individual specifications.
- *Volume flexibility* – changing levels of demand often cause problems, as you can see during the morning rush hour in any major city. Volume flexibility allows an organization to responds quickly to such changes.
- *Diversification or specialization* – which describes how wide a range of products is offered.
- *Technology* – developing and using the latest technologies.

- *Location* – using convenient and cost-effective sites.
- *Time compression* – which is a form of lean operations that concentrates on wasted time.
- *Environmental protection* – a small, but increasing, number of organizations are focusing on sustainable operations and environmental protection.
- *Increased productivity* – using available resources as fully as possible.
- *Adding value* – expanding the product definition to add as much customer value as possible.
- *Growth* – aiming for economies of scale to give both lower costs and better service.

As you can imagine, any particular strategic focus puts different demands on stocks – so the strategic focus clearly influences decisions about stock. But when an organization chooses its focus, it must bear in mind the results that inventory managers can actually deliver. If, for example, stocks are concentrated in a single huge logistics centre, it is unlikely that they could adequately support a focus on flexible, personalized service. Although inventory managers could change their operations, this would be a major undertaking with consequent effects on other activities. So the way that an organization manages its stock can make some strategies difficult or even impossible.

Essentially, then, the strategic decisions must look for the best balance between the organization's internal strengths and the external constraints – matching what the organization is good at to what customers want. This sets the overall framework for inventory decisions, shown in Figure 2.4. Each decision has to work with the organization's logistics strategy (which is set by the higher strategies), its internal strengths (which come from its operations, resources, skills, etc.) and the external environment (which is set by customer requirements, competition, the economic climate, etc.).

Implementing the strategies

The framework for inventory decisions works at all levels. Strategic decisions about stock support the business and logistics strategies, build on long-term internal strengths and consider long-term external factors; tactical decisions about stock support the strategic inventory decisions, and work with medium-term internal and external factors; operational decisions support the tactical inventory decisions, and work with short-term internal and external factors.

We can illustrate this framework by looking at, say, medium-term aims. As a starting point, we might find that the logistics strategy focuses on low costs, so we immediately look for ways of making stock holdings as efficient as possible. We might build on internal strengths to improve the flow of products, reduce the amount of stock, the cost of holding it, the number of times there are shortages, and so on. We cannot, however, only aim at efficient operations, as we have to maintain a level of service that satisfies customers and is competitive. So we might

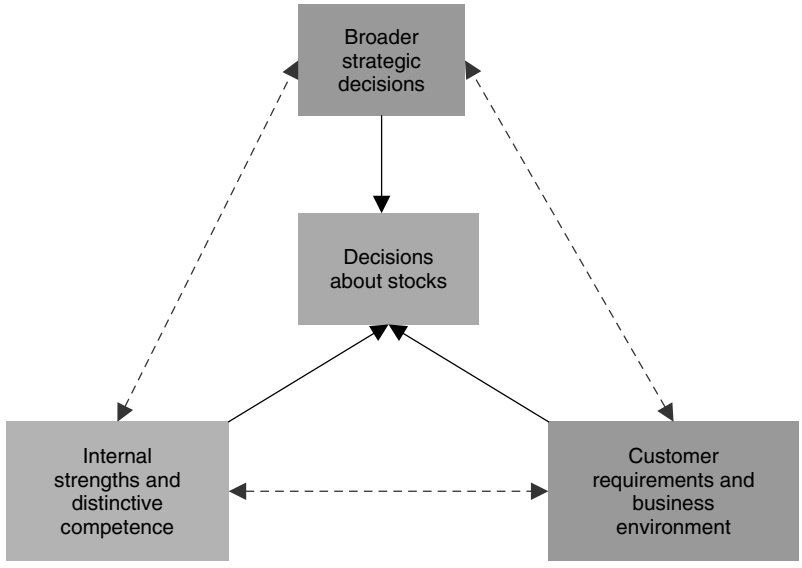

Broader strategic decisions
Decisions about stocks
Internal strengths and distinctive competence
Customer requirements and business environment

Figure 2.4 Framework for inventory decisions

define an immediate aim of achieving some specified level customer service at minimum cost. The level of customer service is a positive decision that is made on the basis of internal and external factors.

Some people describe this balance of service and cost as a *utility*, which is essentially a perceived value. If materials are in the right place, they have *place utility*; if they are available at the right time they have *time utility*. Then we can phrase an aim for stocks in terms of maximizing customer utility. The problems, of course, are getting an accurate measure to describe customer utility, and then finding the decisions that deliver it.

When we describe this approach to inventory management, we obviously make a number of assumptions. One of the basic ones is that inventory managers have the necessary motivation, skills and information to make appropriate decisions. In essence, we assume that they have some basic requirements, including:

- knowledge of the higher levels of decisions and their aims and restrictions;
- understanding of external and internal factors that affect operations;
- ability to define specific aims for stock in terms that can be measured and checked;
- ability to design operations that can achieve these aims;
- analyses for comparing operations and measuring performance – and people who know how to use them;
- ability to manage stocks and actually implement the identified operations;
- accurate data for monitoring and controlling operations.

Summary

Inventory managers have to make decisions at different levels. Their overall aims are set by the organization's broader strategies. Two main strategies are based on lean or agile operations. Under these umbrellas organizations can have a strategic focus on different features, each of which gives different requirements of stocks. This sets a framework for all decision about stocks, which work with higher decisions, external and internal factors.

Review questions

2.3 At what level in an organization are decisions made about stocks?
2.4 Senior logistics managers make all the decisions about inventories. Do you think this is true?
2.5 What is the difference between lean and agile strategies?
2.6 Does the strategic focus have any effect on stocks?
2.7 What are the three main factors that set the context for decisions about stocks?

Strategic role of stock

In the last section we saw how business and logistics strategies set the context for inventory management, and at the same time the actual performance of stock holding can determine which strategies are realistic or even feasible. In this sense, stocks have a clear strategic impact, influencing the long-term options that are available. But we do not have to go this far to see the strategic role of stocks. They have a clear effect on the organization's profit, margins, return on assets, and other financial measures of performance, as well as affecting measures of customer service, such as lead time, availability, perceived product value and reliability. Their ability to decouple production and sales is also a factor in long-term capacity planning, production and productivity.

We could make a long case for the strategic role of stocks and give endless illustrations, but the conclusions are already obvious. Stock has a clear strategic role. We will limit ourselves to two illustrations of this, one allowing smoothed production and a second looking at financial performance.

Stock gives a buffer between production and sales, so that these do not have to match exactly. This gives two considerable benefits. First, smooth operations are much more efficient than variable ones, with easier planning, regular schedules, routine workflow, fewer changes, etc. Second, the organization does not have to install enough capacity to match peak sales – with facilities lying idle and giving low productivity during quieter periods. Any variation between production and sales is covered by changes in stock. When production is higher than sales, stock builds up; when sales are higher than production, stock declines. These changes in stock might be simple adjustments to the finished goods, or they might be changes to work in progress caused by varying the speed of operations. An example of

this comes from Polianski Traders who recorded the following figures over six quarters (values are in millions of Euros).

Quarter	1	2	3	4	5	6
Sales	70	110	120	105	90	80
Percentage change	–	+57	+9	−4	−14	−11
Production	100	105	105	95	95	95
Percentage change	–	+5	0	−10	0	0
Change in stock	+30	−5	−15	−10	+5	+15

As you can see, the stocks allowed production to remain fairly stable – changing by less than 10 per cent – during a period when sales varied by up to 57 per cent. In quarter 3, the company met a record high demand without increasing production, and their costs actually fell as they reduced investment in stock. If the company had installed enough capacity to meet this peak demand, it would have a utilization of only 58 per cent in quarter 1. Proper inventory management allowed higher sales, lower production costs and higher profits.

You can see another long-term effect of stocks in an organization's financial performance. We can illustrate this with the return on assets (ROA), which is defined as the pre-tax profit earned divided by the value of the assets employed.

$$\text{Return on assets} = \frac{\text{profits earned}}{\text{assets employed}}$$

This gives a measure of how well available resources are used and, in general, higher values show better performance. Stocks affect this ratio in several ways (see Figure 2.5).

1. *Current assets*. Assets are conventionally described as current (cash, accounts receivable, stocks, etc.) or fixed (property, plant, equipment, etc.). Better management can reduce stock levels and hence the current assets. Lower investment in stock also frees up cash for more productive purposes and reduces the need for borrowing (see Walters, 1992).
2. *Fixed assets*. Stocks need related investments in warehouses, materials handling equipment, information systems and other facilities which form part of the fixed assets. Reducing stock levels can, therefore, bring associated reductions in fixed assets.
3. *Sales*. Careful management of stocks increases the availability of products, reduces lead times, allows proper delivery size and frequency, and gives faster delivery. This raises their perceived value and gives higher customer satisfaction. The result is more frequent orders from customers, more repeat orders, greater customer loyalty, new customers and generally higher sales.
4. *Profit margin*. More efficient inventory management gives lower operating costs. It can also improve procurement, monitor and control stock levels, set

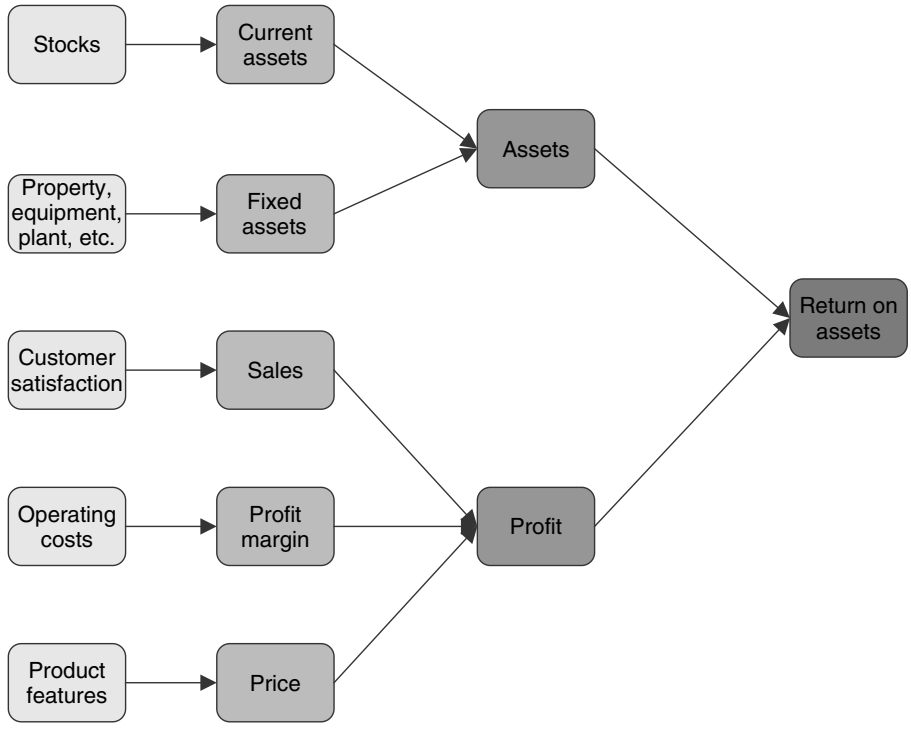


Figure 2.5 Effect of stocks on the return on assets

optimal order sizes and generally reduce inventory costs. The result is higher profit margins, or price reductions to increase sales.

5. *Price*. Point 3 suggested that stocks could raise the perceived value of a product, and then customers are willing to pay a premium price. Stocks can also allow some finishing operations to add value to the overall product package, and again allow an increased price.

The first two points reduce the value of assets, while the last three increase profits. Their combined effect can be a considerable improvement of ROA, with corresponding effects on other measures of performance.

Worked example

CMJ Constructors Ltd. currently has sales of £20 million a year, with a stock level of 25 per cent of sales. Annual holding cost for the stock is 20 per cent of value. Operating costs (excluding the cost of stocks) are £15 million a year and other assets are valued at £30 million. What is the current return on assets? How does this change if stock levels are reduced to 20 per cent of sales?

Solution

Taking costs over a year, the current position is:

Cost of stock	= amount of stock × holding cost	
	= (£20 million × 0.25) × 0.2	= £1 million a year
Total costs	= operating cost + cost of stock	
	= £15 million + £1 million	= £16 million a year
Profit	= sales − total costs	
	= £20 million − £16 million	= £4 million a year
Total assets	= other assets + stock	
	= £30 million + (£20 million × 0.25)	= £35 million
Return on assets	= profit/total assets	
	= £4 million/£35 million	= 0.114 or 11.4%

If stock levels are reduced to 20 per cent of sales:

Cost of stocks	= (£20 million × 0.2) × 0.2	= £0.8 million year
Total costs	= £15 million + £0.8 million	= £15.8 million a year
Profit	= £20 million − £15.8 million	= £4.2 million a year
Total assets	= £30 million + (£20 million × 0.2)	= £34 million
Return on assets	= £4.2 million/£34 million	= 0.124 or 12.4%

Reducing stocks gives lower operating costs, higher profit and a significant increase in the return on assets.

It is clear that management of stocks affects the overall performance of an organization. It can show which strategies are realistic or even feasible; it can increase customer satisfaction and sales; it affects the financial performance; it can increase productivity of resources by using space, capacity, equipment, buildings, capital and staff more efficiently; it affects locations and allows expansion into new market areas. As well as these direct effects, stocks can add value by including additional services, such as providing management information, direct deliveries, vendor-managed inventory, direct deliveries, cross-docking, packaging, postponement, installation and a range of associated services.

The recognition that stocks have a strategic role is a surprisingly recent observation and is one of the most important developments in recent years. It changes the way that organizations manage their supply chains, allowing new processes and encouraging a more integrated view of operations. Any belief that stocks are unused stores of goods sitting idly in warehouses

has been replaced by the view that they are essential supplies of materials moving efficiently through supply chains and providing a crucial service to operations.

Summary

The management of stock has a clear strategic impact on an organization. It influences the feasibility of some operations and directly affects many aspects of organizational performance.

Review questions

2.8 How can stocks make some business strategies infeasible?

2.9 Lower stocks generally give lower costs and better performance. Do you think this is true?

Costs of holding stock

Value of stocks

We have already said that stocks can be very expensive. Organizations obviously look for ways of reducing their overheads as much as possible, but they are often left with surprisingly high costs for stock. Unfortunately, it is difficult to put a precise figure on these. How, for example, do you put a cost on the loss of goodwill when you run out of stock and cannot meet a customer order? There is a good deal of uncertainty in the area and reliance on accounting conventions. We can illustrate this by the problem of valuing stock.

Stocks appear as current assets in a company's accounts. It is, therefore, important to have an accurate value for stocks, as this affects the overall value and performance of the company. In principle, we can find the stock value of any item by simply multiplying the number of units in stock by the unit cost. Unfortunately, the stock level might vary quite widely, so we have the alternatives of taking a snapshot at some specific point, or using an average value over some period. In addition, the stocks are normally bought over some period so the unit cost changes because of inflation, discounts, variations in quality, alternative suppliers, different options or features, and variations in trading conditions and terms. We could also look at the price we expect to get for the stock rather than the amount it costs, but this becomes more speculative. So how can we put a reasonable value on the stock? There are four main methods of doing this (which we shall return to in Chapter 6).

1. *Actual cost*. If an organization holds a few units in stock, it puts an actual cost on each unit. Car showrooms, for example, have a relatively small number of cars in stock, and they can record the actual cost of each unit. This has the

advantage of being accurate and responsive – if the value of a unit changes, this is reported precisely in the stock valuation. But it has the disadvantage of needing to identify and track every unit. It also concentrates on how much each unit cost originally, but not how much it would now cost to replace. For most organizations the turnover of units is too high for actual costs, so they use some method based on average costs.

2. *First-in-first-out (FIFO)* is the standard convention in many countries, and assumes that the units bought earliest are used first. Then the value of stock is set by the amount paid for the last units bought. This approach gives a reasonable figure for the replacement cost, but can – particularly if prices are rising quickly – over-estimate the stock value. There is also no guarantee that customers will pay more for an item whose price has risen. The high values also affects some measures of performance, reducing, say, the return on assets.

3. *Last-in-first-out (LIFO)* is the opposite of FIFO and assumes that the units bought last are sold first. This might be realistic for stocks of coal, for example, where the latest arrival on a tip is used first. LIFO is less widely accepted, and has the disadvantage of assuming that stock always consists of the units bought earliest. If prices are rising, this under-estimates the value of stock. In some circumstances this is an attractive alternative, typically reducing tax liabilities and increasing return on assets.

4. *Weighted average cost*. This finds the average unit cost of all purchases over some time, and assigns this value to all remaining stock. In effect, it adds the total cost of all purchases over a period, and divides this by the total number of units bought, to give a weighted unit cost. It has the advantages of being easy to use, only needing a single calculation in an accounting period (when final accounts are being prepared), and giving a reasonable value of the cost of stock. On the downside, it can under-value stock when costs are rising, as units bought some time ago reduce the average below current values.

The valuation of stock is important in itself, but it also affects other measures, such as the profit margins and gross profits. FIFO assumes that the units sold are the earliest arrivals which usually have lowest cost. This gives higher profit margins than LIFO, which assumes the units sold are the latest, which are likely to cost more.

Worked example

Ulrika Harkness records the following monthly purchases and sales of an item. Assuming she has no opening stock, what is the value of her stock at the end of the period? What is the profit and profit margin if each unit sold for €35?

Month	Number bought	Cost of each unit (€)	Number sold
November	110	22	73
December	60	26	71
January	70	30	49
February	50	28	53
March	80	24	37
April	40	32	71

Solution

Total cost of purchases = $(110 \times 22) + (60 \times 26) + (70 \times 30) + (50 \times 28) + (80 \times 24) + (40 \times 32) =$ €10,680. Stock movements show 410 purchases and 354 sales, so the closing stock is 56 units. Each unit sold for €35, so the income is $35 \times 354 =$ €12,390.

The gross profit is the income minus the cost of units sold:

$$\begin{aligned} \text{Profit} &= \text{income} - \text{cost of units sold} \\ &= \text{income} - (\text{cost of purchases} - \text{cost of remaining stock}) \\ \text{Profit margin} &= \text{profit/number of units sold} \end{aligned}$$

The variable in these calculations is the cost of remaining stock, which we can find using the four conventions described.

1. The *actual cost* needs the price paid for each unit remaining in stock, so we need more information to calculate this. A more detailed analysis of the figures might find that the remaining stock consists of 7 units bought in November, 6 bought in December, 8 bought in January, 11 in February, 17 in March and 7 in April, to give the value of stock as:

$$(7 \times 22) + (6 \times 26) + (8 \times 30) + (11 \times 28) + (17 \times 24) + (7 \times 32) = €1{,}490$$

   Then:

$$\begin{aligned} \text{Profit} &= 12{,}390 - (10{,}680 - 1{,}490) = €3{,}200 \\ \text{Profit margin} &= 3{,}200/354 = €9.04 \text{ a unit} \end{aligned}$$

2. *First-in-first-out* assumes that the 56 units remaining are the last 56 bought, and their value is:

$$(40 \times 32) + (16 \times 24) = €1{,}664$$

   Then:

$$\begin{aligned} \text{Profit} &= 12{,}390 - (10{,}680 - 1{,}664) = €3{,}374 \\ \text{Profit margin} &= 3{,}374/354 = €9.53 \text{ a unit} \end{aligned}$$

3. *Last-in-first-out* assumes that the 56 units remaining are the first 56 bought, and their value is:

$$(56 \times 22) = €1{,}232$$

Then:

$$\begin{aligned}\text{Profit} &= 12{,}390 - (10{,}680 - 1{,}232) = €2{,}942\\ \text{Profit margin} &= 2{,}942/354 = €8.31 \text{ a unit}\end{aligned}$$

A better variation on this calculation says that all 56 units cannot be left over from November, as 110 were bought in the month and 73 were sold, so only 37 could be left. Then in December 60 were bought and 71 were sold, so only 26 of the original delivery are left. In January 70 were delivered and 49 were sold, adding 21 to stock. In February 50 were bought and 53 were sold, so we assume that three were used from January's delivery. Then in March 80 were delivered and only 37 were used, with the remainder going into stock. So our final stock consists of 26 units from November, 18 units from January and 12 units from March, giving a value of:

$$(26 \times 22) + (18 \times 30) + (12 \times 24) = €1{,}400$$

Then:

$$\begin{aligned}\text{Profit} &= 12{,}390 - (10{,}680 - 1{,}400) = €3{,}110\\ \text{Profit margin} &= 3{,}110/354 = €8.79 \text{ a unit}\end{aligned}$$

4. The *weighted average cost* of a unit is:

$$\frac{\text{Total cost of purchases}}{\text{Total number bought}} = \frac{(110 \times 22) + (60 \times 26) + (70 \times 30) + (50 \times 28) + (80 \times 24) + (40 \times 32)}{110 + 60 + 70 + 50 + 80 + 40}$$

$$= 10{,}680/410 = €26.05$$

Then the value of current stock is $(56 \times 26.05) = €1{,}459$

Then:

$$\begin{aligned}\text{Profit} &= 12{,}390 - (10{,}680 - 1{,}459) = €3{,}169\\ \text{Profit margin} &= 3{,}169/354 = €8.95 \text{ a unit}\end{aligned}$$

The actual value of stock is€ 1490; FIFO over-estimates this by 12 per cent, LIFO under-estimates this by 17 per cent or 8 per cent depending on the calculation, and the weighted average is within 2 per cent. The profit varies between€2,942 or€8.31 a unit to€3,374 or€9.53 a unit – a difference of 15 per cent.

Types of cost

All stocks incur costs. These vary widely, but are typically around 20 per cent of the value held a year. It is not surprising that organizations want to minimize their inventory costs but they cannot do this by simply reducing stock. Sometimes low stocks give a minimum cost, but this is not inevitable and low stocks can lead to shortages that disrupt operations and have very high costs. To look at this balance more closely, we need to describe some details of the costs involved. The usual approach describes four types of cost – unit, reorder, holding and shortage costs:

1. *Unit cost*. This is the price charged by suppliers for one unit of the item, or the cost to the organization of acquiring one unit. In general, it is fairly easy to find values for the unit cost by looking at quotations or recent invoices from suppliers. As we have just seen, though, it can be difficult to find an accurate unit cost, and this is particularly true when there are several suppliers offering alternative products or giving different purchase conditions. If the company makes the item itself, it can be difficult to set a reliable production cost or to calculate a valid transfer price.

2. *Reorder cost*. This is the cost of placing a repeat order for the item and might include allowances for drawing-up an order (with checking, getting authorization, clearance and distribution), correspondence and telephone costs, receiving (with unloading, checking and testing), supervision, use of equipment and follow-up. Sometimes costs such as quality control, transport, delivery, sorting and movement of received goods are included in the reorder cost.

   The reorder cost should be the cost of repeat orders and not first-time purchases, which might include allowances for finding suitable suppliers, checking their reliability and quality, requesting quotations, negotiations with alternative suppliers, and so on. In practice, you might find the best estimate for a reorder cost by dividing the total annual cost of the purchasing department (plus any other relevant costs) by the number of orders sent out.

   A special case of the reorder cost occurs when an item is made within the organization. Then the reorder cost is typically a batch set-up cost which might include production documentation, production lost while resetting machines, idle time for operators, material spoilt in test runs, slowed production during familiarization, and so on.

3. *Holding cost*. This is the cost of holding one unit of an item in stock for one period of time. As the usual period for calculating stock costs is a year, a holding cost might be expressed as, say, £10 a unit a year.

   The most obvious cost of holding stock is money tied up – which is either borrowed (in which case there is interest to pay), or could be put to other use (in which case there are opportunity costs). Other holding costs are due to storage space (supplying a warehouse, rent, rates, heat, light, etc.), loss (due to damage, obsolescence and pilferage), handling (including all movement, special packaging, refrigeration, putting on pallets, etc.), administration (stock checks, computer updates, etc.) and insurance.

It is difficult to suggest values for these, but one view has percentages of unit cost, as:

	% of unit cost
cost of money	10–15
storage space	2–5
loss	4–6
handling	1–2
administration	1–2
insurance	1–5
Total	19–35

4. *Shortage cost*. If an organization runs out of stock for an item and there is demand from a customer, then there is a shortage that has an associated cost. In the simplest case a retailer might lose the profit from a lost sale. Usually, though, the effects of shortages are wider than this and include loss of goodwill, loss of future sales, loss of reputation, and so on. Any shortage of parts for production might cause considerable disruption and force emergency procedures, rescheduling of operations, retiming of maintenance period, laying-off employees, and so on. Shortage costs might also contain allowances for positive action to counteract the shortage, such as sending out emergency orders, paying for special deliveries, using alternative and more expensive suppliers, or storing partly finished goods.

   Many of these shortage costs are difficult to measure and are often little more than informed guesses. We might estimate the cost of lost future sales, for example, but there is no way of getting an exact figure. It is often suggested that shortage costs are inherently inaccurate and can be misleading, so they should only be used with caution. There is, however, general agreement that shortages can be damaging, and organizations are often willing to incur the relatively lower costs of carrying stock to avoid potential shortages.

Stock turnover

The costs obviously vary widely from organization to organization, but a typical value is around 20 per cent of value held a year. Whether this figure is rising or falling is open to debate. Some people say that land, building, fuel, taxes, safety, environmental protection and employee costs are all rising and making stocks more expensive. An opposing view says that improvements in inventory management are more than compensating for any rising costs, and the overall result is a decline. The real picture depends on circumstances within each organization.

It might be difficult to use costs to monitor inventory performance over time, so we can use some more direct measures. Perhaps the most common measure

is the stock turnover, which is the ratio of the number of units sold to the average stock:

$$\text{Stock turnover} = \frac{\text{number of units sold in a period}}{\text{average stock}}$$

Car assembly plants and oil refineries achieve turnovers of over 50 times a year, with units kept in stock for an average of about a week; department stores have turnovers below 5. If we can increase the stock turnover without affecting customer service, overall costs decrease. Organizations can often make useful comparisons with competitors having similar operations – if our shop has a stock turnover of 20 times a year, we are doing better than the shop next door that only achieves 15 times a year. However, differences in operations, product packages, objectives and other circumstances mean that you should view these comparisons with caution.

Organizations often find their stock turnover indirectly from accounts rather than directly from actual counts. Then we can define a cost of units sold as the total cost of buying or acquiring the units that are later sold to customers, and the turnover is:

$$\text{Stock turnover} = \frac{\text{cost of units sold}}{\text{average value of stock}}$$

This obviously introduces some of the problems with costs and accounting conventions that we mentioned earlier.

Worked example

Emergent Technologies Wholesale (Scandinavia) buys an item for €100 a unit and sells it for €150. Annual sales of the item are around 1,000 units, with average stock of 150 units. Each unit held in stock costs approximately 25 per cent of cost a year.

1. Describe the stock holdings.
2. What are the benefits if average stocks of the item are reduced to 100 units without affecting customer service?

Solution

1. We can calculate a number of measures for the stock.
   - Stock turnover is: $\dfrac{\text{cost of units sold}}{\text{average value stock}} = \dfrac{1000 \times 100}{150 \times 100} = 6.67$
   - Annual gross profit on the item: sales × (selling price − unit cost)
     $= 1{,}000 \times (150 - 100) = €50{,}000$ a year
   - Average investment in stock: number of units held × unit cost
     $= 150 \times 100 = €15{,}000$

- Annual stock holding cost: average number of units in stock × cost of holding each unit = 150 × (100 × 0.25) = €3,750 a year
- Stock holding cost per unit sold: $\frac{\text{stock holding cost}}{\text{number of units sold}} = \frac{3750}{1000} = £3.75$

2. If the average stock is reduced to 100 units, the company would be operating more efficiently with:
   - stock turnover = (1, 000 × 100)/(100 × 100) = 10
   - average investment in stock = 100 × 100 = €10,000
   - annual stock holding cost = 100 × (100 × 0.25) = €2,500
   - stock holding cost per unit sold = 2,500/1,000 = €2.50
   - additional profit = reduced stock holding cost = 3,750 − 2,500 = €1,250 a year

Summary

Holding stocks is expensive, typically costing around 20 per cent of value a year. This cost is formed from the four components of unit, reorder, holding and shortage costs. It can be difficult to find precise values for these. Even the basic figure of stock value depends on the conventions used.

Review questions

2.10 To what extent does the value of stock depend on the accounting conventions used?

2.11 Very roughly, what is the cost of storing a freezer full of food in a kitchen?

2.12 Should organizations aim for a high stock turnover or a low one?

2.13 What is the main purpose of measuring inventory performance?

Approaches to inventory control

Basic questions

Inventory managers make a series of decisions about stock such as where do we hold stock and in what kind of facilities? What suppliers and transport operators do we use? What information systems do we use? Can we form alliances? These, and all the other decisions about stocks, set the overall context. There comes a point, however, when inventory managers have to make some immediate decisions about their stocks. Having set the scene with a series of appropriate decisions, they are now responsible for some actual stocks, and have to start looking after them. What decisions do they make?

As always, there is no shortage of possible topics, but we will concentrate on the basic questions of inventory control. In particular, we will look at three fundamental questions; 'What items should we keep in stock?', 'When should we

place an order with suppliers?' and 'How much should we order?'. Answers to these three questions set most of the features of the stocks.

1. *What items should we keep in stock?* Holding any stock is expensive, so organizations have to make sure that their stocks remain at the lowest level that allows acceptable service. This means:

   - keeping stock of existing items at reasonable levels;
   - not adding unnecessary items to the inventory;
   - removing all items which are no longer used from the inventory.

   Unless tightly controlled, there is a tendency for stock holdings to drift upwards. In organizations that are not particularly short of money or storage space, this rise can be quite fast. In part, this is a result of increasing stocks of existing items, based on arguments like, 'we are succeeding because we give good customer service and this needs higher stocks'. In part, it is a consequence of continually adding new items to the inventory. As an organization's operations evolve, its requirements for stocks change, and it adds new items to replace older ones that it no longer needs. In reality, new items are often introduced without much planning, while old ones are left in stock on the off-chance that they are needed again. This is particularly common with inventories of spare parts. When a new machine is bought to replace an older one, the new spare parts are added to stock, but nobody remembers to remove the spares for the replaced machine. It is often difficult to accept that an item that was expensive to buy now has no value and should simply be scrapped.

   If stocks continue to grow, they eventually cause concern, perhaps when supplies of space or money become scarcer. Then managers look for some rationalization. Unfortunately, they often want fast results and are tempted to give instructions like, 'stock holdings must fall by 10 per cent'. The easiest way of achieving this is by not replacing items when they are issued. As a result, fast-moving items that are used frequently have their stock levels reduced, while slow-moving items that are seldom, if ever, used are left on the shelves at their original levels.

   Of course, these situations should never occur, as stocks should be controlled by rational policies that relate holdings to demand. In particular, before any item is added to their inventory, managers should compare the costs and benefits of holding it, to those of not holding it. Then they should only add an item if this gives clear benefits. In the same way, managers should monitor the use of items already on the inventory, and when it becomes cheaper to no longer stock them, they should be removed as quickly as possible.

2. *When should we place an order?* There are basically three different approaches to this question. The first uses a periodic review to place orders of variable size at regular intervals of time. Any variation in demand is allowed for by the changing order size. This system is often used in shops, where stocks are reviewed at the end of a day and any units sold are replaced.

   The second approach uses a fixed order quantity. Stock levels are continuously monitored and when they fall to a specified level a fixed amount is

ordered. Any variation in demand is allowed for by changing the time between orders. A central heating plant, for example, may order 25,000 litres of oil whenever the amount in the tank falls to 2,500 litres.

The third approach relates the supply more directly to the demand and orders enough stock to meet known demand over a specified period. Then both the time and quantity ordered depend directly on demand.

3. *How much should we order?* Every time we place an order, there are associated costs for administration, delivery, and so on. If we place large, infrequent orders, the costs of ordering and delivery are kept low, but stock levels and average inventory value are high. If we place small frequent orders, costs of ordering and delivery are high, but average stock level is low. In general, we will look for a compromise between these two extremes that minimizes overall cost.

We have suggested three basic questions of inventory management, but it might seem that we have omitted several others such as, 'What should average stock levels be?', 'How much should we invest in stock?', 'What service level do we offer?' and 'How often will there be shortages?'. In practice, the answers to these other questions come from the three basic questions. Once we make decisions about the frequency and size of orders, we automatically set the average stock levels, costs, service level, likelihood of shortage and a whole range of other features. We will start to explore this theme in the next chapters.

Answering the basic questions

Having posed the three basic questions, we now have to look for answers. There are basically two different methods, which are based on the ways of assessing demand.

1. *Independent demand* methods assume that the demand for an item is independent of the demand for any other item. Then the aggregate demand for an item is made up of many independent demands from separate customers – in the way that the overall demand for milk in a supermarket is made up of many small demands from separate customers. In these circumstances the only reasonable way of forecasting future demand is to project historical trends. Stock control is then based on quantitative models that relate forecast demand, costs and other variables, to find optimal values for order quantities and timing. These models can be either fixed order quantity or periodic review. We will look at independent demand methods in Part II of this book (see Figure 2.6).
2. *Dependent demand methods*. In practice, the demands for different items are often related: in a café the demand for chips is related to the demand for fish; in a factory the demand for all components of a product are related through the production plans; demand for both domestic electricity and gas are linked through the weather, and so on. In such circumstances we can often find an alternative way of forecasting the demand for an item, and can use this in a dependent demand method.

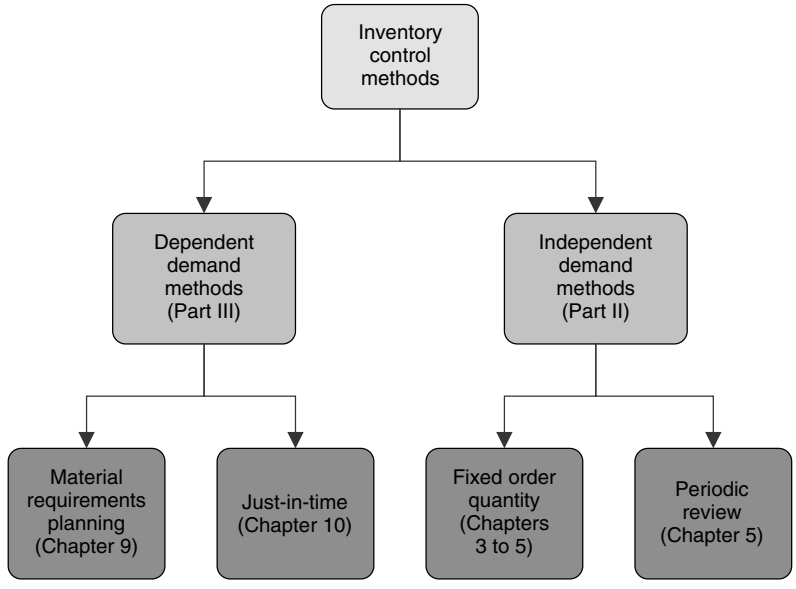

Inventory control methods
Dependent demand methods (Part III)
Independent demand methods (Part II)
Material requirements planning (Chapter 9)
Just-in-time (Chapter 10)
Fixed order quantity (Chapters 3 to 5)
Periodic review (Chapter 5)

Figure 2.6 Types of inventory control system

Common approaches of this type link the demand for components to production plans for finished goods. These methods are formalized into methods like Material Requirements Planning and Just-in-Time which we describe in Part III of this book.

Summary

Inventory control is based on the answers to three basic questions. These ask about the items to stock, the timing of orders and the quantities ordered. There are two approaches to answering these questions. Independent demand methods use forecasts of demand to define optimal order quantities and times; dependent demand methods look for other ways of co-ordinating the supply and demand.

Review questions

2.14 Why is the average stock level not included as a fundamental question of inventory control?

2.15 What three methods are used to set the timing of orders?

2.16 How is demand for an item found for independent demand inventory methods?

Chapter review

- The last chapter introduced the general ideas of inventory management. This chapter focused more specifically on the role of stocks within an organization.

- Inventory management can be viewed as a part of the broader logistics function. Decisions about stocks are linked, in some way, to all other areas of logistics, so it is best to take a broad view and look for the best overall results.
- Inventory managers make decisions at different levels. Their overall aims are set by the organization's broader strategies, with two main strategies based on lean or agile operations. Under these umbrellas organizations can have a strategic focus on different features, each of which gives different requirements of stocks.
- The framework for all decision about stocks is set by higher decisions, and a balance of external and internal factors.
- The management of stock has a clear strategic impact. It influences the feasibility of certain operations and directly affects organizational performance.
- Holding stocks is expensive, typically costing around 20 per cent of value a year. It can be difficult to find precise values for this, and even the basic figure of stock value depends on the conventions used.
- Inventory control is based on the answers to three basic questions. These ask about the items to stock, the timing of orders and the quantities ordered.
- There are two approaches to answering these questions. Independent demand methods use forecasts of demand to define optimal order quantities and times; dependent demand methods look for other ways of co-ordinating the supply and demand.

Project

Many companies report their stock holdings in annual accounts or on their websites. You can see from the following examples that amount of stock (measured as a proportion of sales) varies widely.

- Tesco (2003) the UK's largest retailer with:
  - Sales for first 6 months £12,733 million
  - Stocks £957 million
- ICI (2003) specialist international chemical manufacturer
  - Turnover for first six months £2,821 million
  - Stocks £748 million
- FGH Horsveldt (2003) printer of specialist materials
  - Turnover for first six months €52 million
  - Stocks €35 million
- Persimmon (2003) one of the UK's largest house builders
  - Turnover for first six months £820 million
  - Stocks £1,377 million

The aim of this project is to look for patterns in the stock holdings of different types of company. For this you should take a representative sample of companies from different industries and see if you can find any patterns. Explain why some industries hold more stock than others. What features tend to give high stocks?

Discussion questions

2.1 Why is it better to view inventory management in the broader context of logistics, than to consider it as a separate function?

2.2 Strategic decisions within an organization tend to be vague and unfocused. How can these hazy aspirations have any impact on a precise area like inventory management? How can different business strategies affect the decisions made by inventory managers?

2.3 The aim of inventory management can be summarized as ''achieving a specified level of customer service at minimum cost''. To what extent do you think that this is true?

2.4 Is it really possible for stores of goods sitting in a warehouse to have a strategic impact on an organization?

2.5 The framework for all inventory decisions is set by higher decisions, and a balance of internal and external factors. What exactly does this mean, and is it necessarily true?

2.6 In practice it can be difficult to find the costs of holding stock. Why is this? How can these difficulties be overcome? Are there any other views of inventory costs that could be easier to work with?

2.7 It is a gross simplification to say that there are only three basic questions of inventory control. Do you agree with this?

References and further reading

Beesley, A. (2000) Time compression in the supply chain, Chapter 11 in D. Waters (ed.), *Global Logistics and Distribution Planning*, London: Kogan Page.

Christopher, M. (1996) Emerging issues in supply chain management, *Proceedings of the Logistics Academic Network Inaugural Workshop*, Warwick.

FGH Horsveldt (2003) *Interim Report for 2002*, Copenhagen: FGH Horsveldt.

ICI (2003) *Interim Report for 2002*, London: ICI plc.

Landvater, D.V. (1997) *World Class Production and Inventory Management*, New York: John Wiley.

Monden, Y. (1983) *Toyota Production System*, Atlanta, Georgia: Industrial Engineering and Management Press.

Ohno, T (1988) *Toyota Production System*, New York: Productivity Press.

Persimmon (2003) *Interim Report for 2002*, York: Persimmon plc.

Porter, M.E. (1985) *Competitive Advantage*, New York: Free Press.

Tesco (2003) *Interim Report for 2002*, Cheshunt, Hertfordshire: Tesco plc.

Townsend, R. (1970) *Up the Organisation*, London: Coronet Books.
Walters, D. (1992) Strategic inventory management, *Focus on Logistics and Distribution Management*, 11(5), 4–12.
Waters, D. (2003) *Logistics: An Introduction to Supply Chain Management*, Basingstoke: Palgrave Macmillan.
Womack, J., Jones, D. and Roos, D. (1990) *The Machine that Changed the World*, New York: Rawson.

Part II

Methods for Independent Demand

3

Economic Order Quantity

Aims of the chapter

This chapter introduces some quantitative models for inventory control. The first model takes an idealized stock and finds the fixed order size that minimizes costs. This is the *economic order quantity*, which is the basis of most independent demand methods. Related calculations and extensions to this basic model are developed in the following chapters.

After reading this chapter you should be able to do the following:

- discuss the reasoning behind the economic order quantity;
- derive an equation for the economic order quantity;
- calculate the economic order quantity for an item;
- measure the effects of moving away from the economic order quantity;
- work with integer order sizes;
- measure the effects of errors and approximations in costs and forecast demand;
- calculate a reorder level;
- outline some limitations of the reorder level.

This chapter emphasizes:

- *quantitative models* for inventory control;
- calculation of an *economic order quantity* for the order size;
- calculation of a *reorder level* to find the timing of orders.

Defining the economic order quantity

Background to the model

This chapter describes one of the standard analyses of inventory control. It shows how we can balance the various costs of stock to answer the question, 'How much

should we order?' The approach is to build a model of an idealized inventory system and calculate the fixed order quantity that minimizes total costs. This optimal order size is called the *economic order quantity* (EOQ).

The EOQ calculation is the most important analysis of inventory control, and arguably one of the most important results derived in any area of operations management. The first reference to the work is by Harris (1915), but the calculation is often credited to Wilson (1934) who independently duplicated the work and marketed the results.

The stock level of an item varies over time, with a typical pattern shown in Figure 3.1. At some point, A, a delivery arrives and raises the stock level. Then the stock declines as units are removed to meet customer demands. At some point, B, an order for replenishment is placed with a supplier and this arrives at time C. This general pattern, with some short-term variations, is repeated as long as the item is kept in stock. Sometimes an unexpectedly heavy demand or delayed

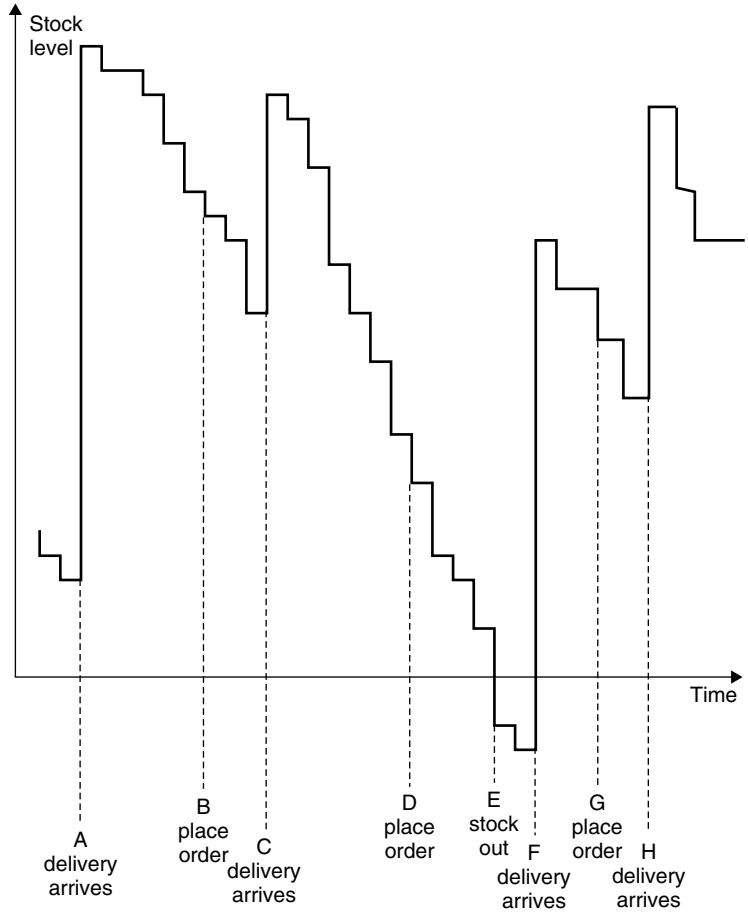


Figure 3.1 Typical pattern of stock level over time

delivery means that stocks run out (as at point E) and then we can represent shortages by negative stock levels. At other times an unexpectedly low demand or fast delivery means that deliveries arrive when they are not really needed (as at points C and H).

We can analyse this pattern, but it is fairly complicated so we start with a basic model that makes a number of assumptions:

- the demand is known exactly, is continuous and is constant over time;
- all costs are known exactly and do not vary;
- no shortages are allowed;
- lead time is zero – so a delivery is made as soon as the order is placed.

A number of other assumptions are implicit in the model, including:

- we can consider a single item in isolation, so we cannot save money by substituting other items or grouping several items into a single order;
- purchase price and reorder costs do not vary with the quantity ordered;
- a single delivery is made for each order;
- replenishment is instantaneous, so that all of an order arrives in stock at the same time and can be used immediately.

Perhaps the most important assumption here is that demand is known exactly, is continuous and constant over time (as shown in Figure 3.2). This, and the other assumptions, might seem unrealistic, but you should remember two things. First, all models are simplifications of reality and their aim is to give useful results rather than be exact representations of actual circumstances. The economic order quantity is widely used, and we can infer that it is accurate enough for many purposes. The results may not be optimal in the strict mathematical sense, but they are good approximations and do, at worst, give useful guidelines. Second, this is a basic model that we can extend in many ways. In the next two chapters we remove some of the assumptions and develop more complex models.

The assumptions give an idealized pattern for a stock level. A continuous demand means that the stock level declines steadily rather than falls in a series of steps; a constant demand means that the decline is always at the same rate. If the lead time is zero, we need never place an order before stock actually runs out – as placing an order when there is stock remaining would leave a residue that is never used and only incurs holding costs. The assumption that no shortages are allowed means that the stock level never falls below zero, and there are no lost sales. Finally, we are looking for the fixed order quantity that will minimize costs, so we always place orders of exactly this size. The resulting pattern is shown in Figure 3.3.

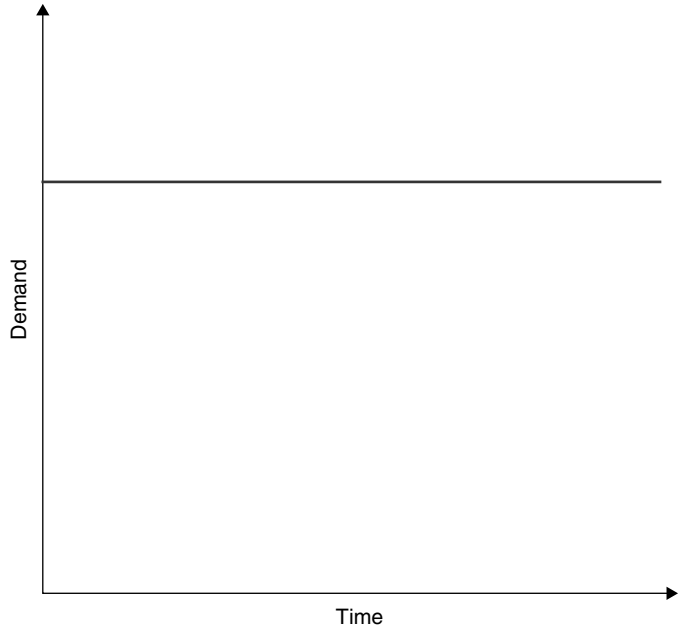


Figure 3.2 Demand is constant and continuous over time

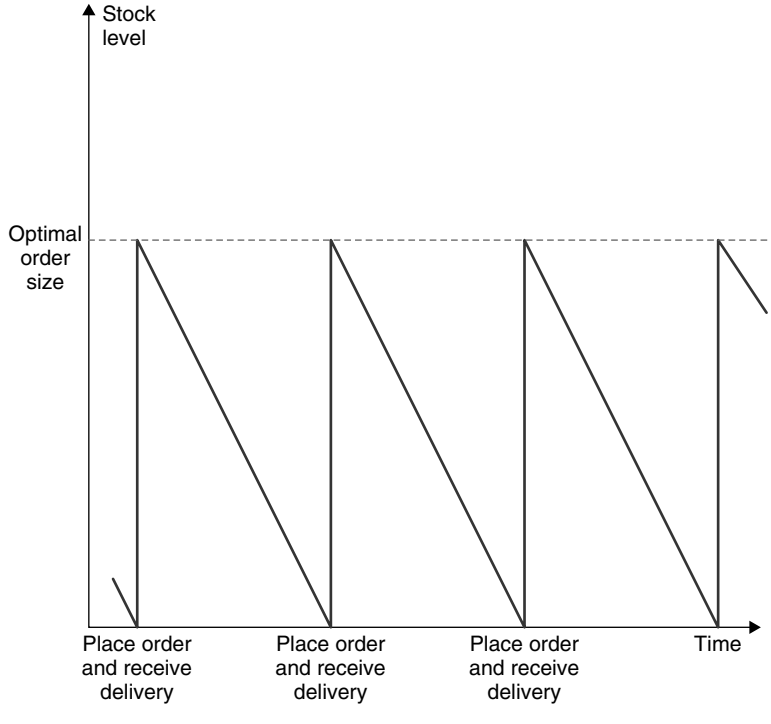


Figure 3.3 Stock level with fixed order size

Variables used in the analysis

Now we have laid the foundations for our model and can introduce some details, starting with a list of variables. In the last chapter we described the four costs of inventory:

- *Unit cost (UC)* is the price charged by the suppliers for one unit of the item, or the total cost to the organization of acquiring one unit.
- *Reorder cost (RC)* is the cost of placing a routine order for the item and might include allowances for drawing-up an order, correspondence, telephone costs, receiving, use of equipment, expediting, delivery, quality checks, and so on. If the item is made internally, this might be a set-up cost.
- *Holding cost (HC)* is the cost of holding one unit of the item in stock for one period of time. The usual period for calculating stock costs is a year, so a holding cost might be, say, £10 a unit a year.
- *Shortage cost (SC)* is the cost of having a shortage and not being able to meet demand from stock. In this analysis we have said that no shortages are allowed, so SC does not appear (it is effectively so large that any shortage would be prohibitively expensive).

If you look at the stock pattern in Figure 3.3, you can see that there are three other variables:

- *Order quantity (Q)* which is the fixed order size that we always use. The purpose of this analysis is to find an optimal value for this order quantity.
- *Cycle time (T)* which is the time between two consecutive replenishments. This depends on the order quantity, with larger orders leading to longer cycle times.
- *Demand (D)* which sets the number of units to be supplied from stock in a given time period (for example, ten units a week). Here, we assume that the demand is continuous and constant.

The only variable that is directly under our control is the order quantity, and we can give this any value we like. When we set the order quantity, this fixes the cycle length. We assume that all the other parameters are fixed and beyond our control. Our aim, then, is to find optimal values for Q – and thereby T – in terms of these other constants.

Derivation of the economic order quantity

This derivation uses a standard approach that is suitable for many stock control models. It has three steps, as follows:

1. Find the total cost of one stock cycle.
2. Divide this total cost by the cycle length to get a cost per unit time.
3. Minimize this cost per unit time.

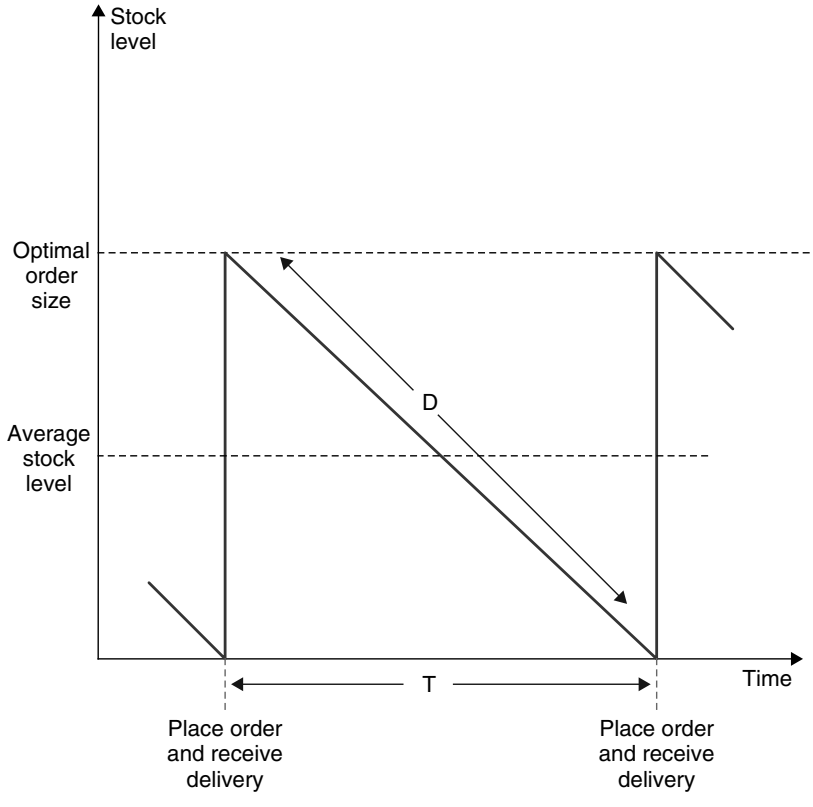


Figure 3.4 Features of one stock cycle

If we take one stock cycle from Figure 3.3, we get the pattern shown in Figure 3.4. At some point we place an order for a quantity, Q, which arrives instantly and is used at a constant rate, D. Eventually no stock remains and it is time to place another order. The cycle has a length T. We know that during the cycle the amount entering stock is Q, while the amount leaving is $D \times T$. These must be equal as the stock level at both the start and finish of the cycle is zero.

$$\text{amount entering stock in cycle} = \text{amount leaving stock in cycle}$$

so

$$Q = D \times T$$

The first step of the analysis finds the total costs for a cycle, and we find this by adding the three separate components of cost for units, reorders, and holdings (remembering there are no shortage costs). Hence:

$$\begin{array}{lllllll} \text{total cost per cycle} & = & \text{unit cost component} & + & \text{reorder cost component} & + & \text{holding cost component} \end{array}$$

We can calculate these separate components for a cycle as follows:

- unit cost component = unit cost (UC) × number of units ordered (Q)
  $= UC \times Q$
- reorder cost component = reorder cost (RC) × number of orders placed (1)
  $= RC$
- holding cost component = holding cost (HC) × average stock level (Q/2) × time held (T)
  $= \frac{HC \times Q \times T}{2}$

Adding these three components gives the total cost per cycle as:

$$\text{Total cost per cycle} = UC \times Q + RC + \frac{HC \times Q \times T}{2}$$

This completes the first step of the analysis. The second step divides this cost by the cycle length, T, to give a total cost per unit time, TC:

$$\text{Total cost per unit time} = TC = \frac{UC \times Q}{T} + \frac{RC}{T} + \frac{HC \times Q}{2}$$

But we know that $Q = D \times T$ or $D = Q/T$ and substituting this gives:

$$TC = UC \times D + \frac{RC \times D}{Q} + \frac{HC \times Q}{2}$$

The demand and all costs are fixed, so the only variable on the right-hand side of this equation is Q. So we can see how the total cost per unit time varies with order quantity. The most convenient way of doing this is to plot each of the cost components separately against Q and then add them together (as shown in Figure 3.5).

The unit cost component (UC × D) is independent of order quantity and is 'fixed'. The other two components vary with order quantity and form the 'variable' cost per unit time. In particular, the holding cost component rises linearly with Q while the reorder cost component falls as Q increases. Adding the three components together gives a total cost curve which is an asymmetric 'U' shape with a distinct minimum. This minimum corresponds to the optimal order size. With orders smaller orders than this, costs rise because of the higher reorder cost component; with orders larger than this, costs rise because of the higher holding cost component.

The third step of our analysis finds the minimum cost per unit time. For this we differentiate the equation for TC with respect to Q and set the result equal to zero:

$$\frac{d(TC)}{dQ} = -\frac{RC \times D}{Q^2} + \frac{HC}{2} = 0$$

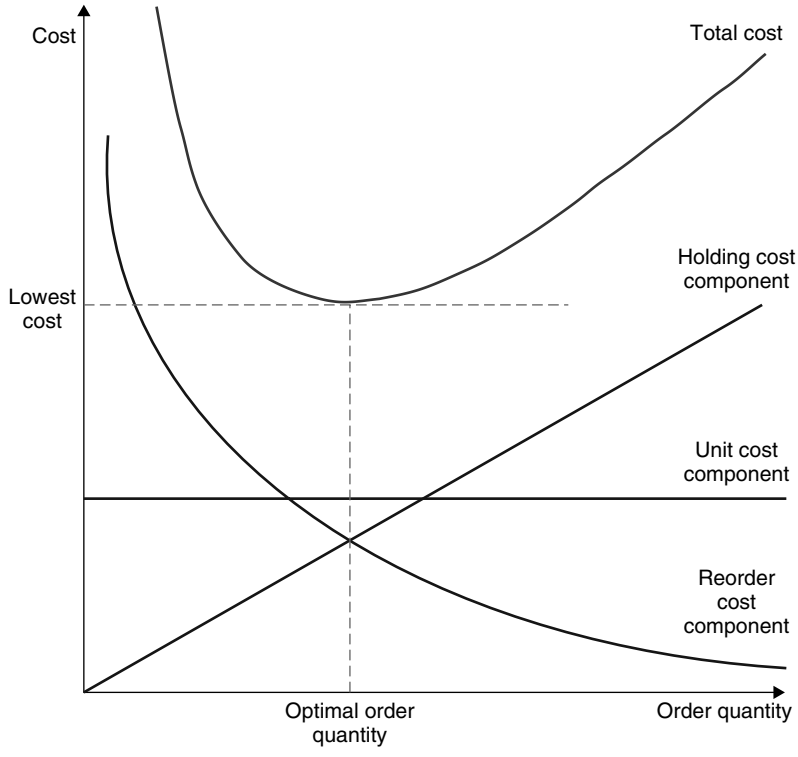


Figure 3.5 Variation of costs with order quantity

When we rearrange this we get the optimal order size, or the economic order quantity, which we will call Qo:

$$\text{Economic order quantity} = \text{Qo} = \sqrt{\frac{2 \times \text{RC} \times \text{D}}{\text{HC}}}$$

This is the most important result of the analysis and answers the question, 'How much should we order?' Now we can find the corresponding optimal length of the stock cycle. We know that Q = D×T, and if we substitute Qo for Q we find the optimal cycle length, To.

$$\text{optimal cycle length} = \text{To} = \text{Qo/D} = \sqrt{\frac{2 \times \text{RC}}{\text{D} \times \text{HC}}}$$

We can also find the optimal cost per unit time, TCo, by substituting the value for Qo. We know that:

$$\text{TC} = \text{UC} \times \text{D} + \frac{\text{RC} \times \text{D}}{\text{Q}} + \frac{\text{HC} \times \text{Q}}{2}$$

The unit cost component is fixed, so we can concentrate on the last two terms which form the variable cost (VC). Then:

$$VC = \frac{RC \times D}{Q} + \frac{HC \times Q}{2}$$

Substituting for Qo to give the optimal value VCo:

$$\begin{aligned} VCo &= RC \times D \times \sqrt{\frac{HC}{2 \times RC \times D}} + \frac{HC}{2} \times \sqrt{\frac{2 \times RC \times D}{HC}} \\ &= \sqrt{\frac{RC \times HC \times D}{2}} + \sqrt{\frac{RC \times HC \times D}{2}} \\ &= \sqrt{2 \times RC \times HC \times D} \end{aligned}$$

If you compare this to the economic order quantity, you can see that:

$$\boxed{\text{Optimal variable cost per unit time} = VCo = HC \times Qo}$$

Then the optimal total cost per unit time is the sum of this variable cost and the fixed cost:

$$\boxed{\text{optimal cost per unit time} = TCo = UC \times D + VCo}$$

You can see an interesting point in the equation above which clearly shows that for the economic order quantity the reorder cost component equals the holding cost component (both have the value $\sqrt{RC \times HC \times D/2}$).

Worked example

Jaydeep (Trading) Company buys 6,000 units of an item every year with a unit cost of $30. It costs $125 to process an order and arrange delivery, while interest and storage costs amount to $6 a year for each unit held. What is the best ordering policy for the item?

Solution

Listing the values we know in consistent units:

Demand	= D	=	6,000 units a year
Unit cost	= UC	=	$30 a unit
Reorder cost	= RC	=	$125 an order
Holding cost	= HC	=	$6 a unit a year

Substituting these figures into the economic order quantity equation gives:

$$Qo = \sqrt{\frac{2 \times RC \times D}{HC}} = \sqrt{\frac{2 \times 125 \times 6{,}000}{6}} = 500 \text{ units}$$

The optimal time between orders is:

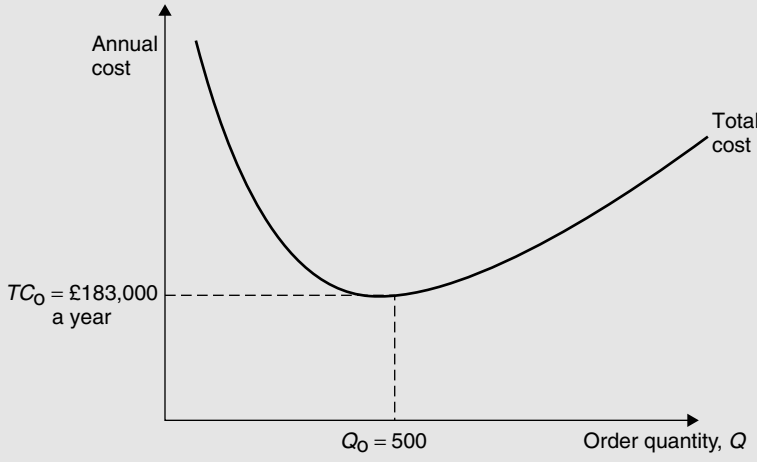
$$\text{To} = \text{Qo}/\text{D} = 500/6{,}000 = 0.083 \text{ years} = 1 \text{ month}$$

The associated variable cost is:

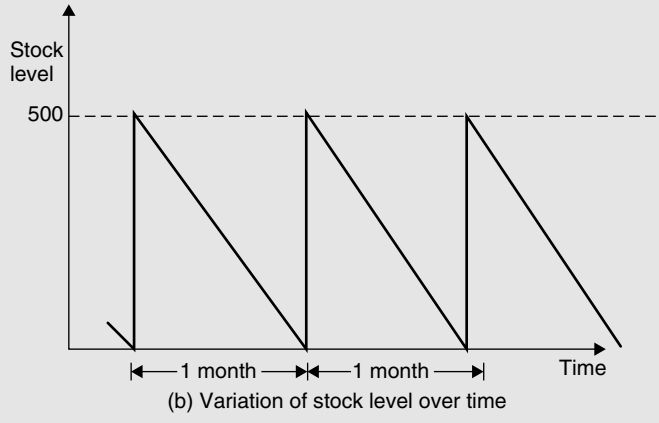
$$\text{VCo} = \text{HC} \times \text{Qo} = 6 \times 500 = \$3{,}000 \text{ a year}$$

This gives a total cost of:

$$\text{TCo} = \text{UC} \times \text{D} + \text{VCo} = 30 \times 6{,}000 + 3{,}000 = \$183{,}000 \text{ a year}$$



(a) Variation of total cost per unit time with order size



(b) Variation of stock level over time

Figure 3.6 Variations in total cost and stock level

The optimal policy is to order 500 units a month, with annual costs of $183,000 (shown in Figure 3.6).

Worked example

Sarah Brown works for a manufacturer that makes parts for marine engines. The parts are made in batches, and every time a new batch is started it costs £1,640 for disruption and lost production and £280 in wages for the fitters. One item has an annual demand of 1,250 units with a selling price of £300, 60 per cent of which is direct material and production costs. If the company looks for a return of 20 per cent a year on capital, what is the optimal batch size for the item and the associated costs?

Solution

From the information given:

- The annual demand (D) is 1,250 units.
- As 60 per cent of the selling price is made up of direct costs, the unit cost (UC) becomes $0.6 \times 300 = £180$.
- Annual holding cost (HC) is 20 per cent of unit cost, or $0.2 \times 180 = £36$.
- Reorder, or in this case batch set-up, cost (RC) has two components for lost production and direct wages and becomes $1{,}640 + 280 = £1{,}920$.

Substituting these gives the optimal order size:

$$Qo = \sqrt{\frac{2 \times RC \times D}{HC}} = \sqrt{\frac{2 \times 1{,}920 \times 1{,}250}{36}} = 365 \text{ units (about)}$$

The optimal time between orders is:

$$To = Qo/D = 365/1{,}250 = 0.29 \text{ years} = 15 \text{ weeks (about)}$$

The optimal variable cost is:

$$VCo = HC \times Qo = 36 \times 365 = £13{,}140 \text{ a year}$$

The optimal total cost is:

$$TCo = UC \times D + VCo = 180 \times 1{,}250 + 13{,}140 = £238{,}140 \text{ a year}$$

Doing the calculations

In practice, a huge amount of software has been developed for inventory control. Much of this uses specialized programs, which can be very large and complicated. You can, however, find software that does some of the analysis, or illustrates the principles. Spreadsheets can be useful here – as illustrated in Figure 3.7, which shows a basic calculation of the economic order quantity. Information is entered

	A	B	C	D	E	F	G
1	**Inventory Control - Economic Order Quantity**						
2							
3	**Note:**	Using the economic order quantity:					
4		order size, Q = sqrt((2*RC*D)/HC)			where:	UC = unit cost	
5		fixed unit costs = UC*D				RC = reorder cost	
6		variable cost = RC*D/Q + HC*Q/2				HC = holding cost	
7						D = demand	
8						Q = order quantity	
9							
10	**Inputs**	Demand		1,200			
11		Unit cost		£ 25.00			
12		Holding cost		£ 7.50			
13		Reorder cost		£ 50.00			
14							
15	**Results**	Order size		126.49			
16		Fixed unit costs		£ 30,000.00			
17		Variable cost		£ 948.68			
18		Total cost a period		£30,948.68			
19							
20	**Costs**	**Order size**	**Reorder cost**	**Holding cost**	**Variable cost**		
21		51	£ 1,186	£ 190	£ 1,376		
22		76	£ 791	£ 285	£ 1,075		
23		101	£ 593	£ 379	£ 972		
24		126	£ 474	£ 474	£ 949	minimum	
25		152	£ 395	£ 569	£ 964		
26		177	£ 339	£ 664	£ 1,003		
27		202	£ 296	£ 759	£ 1,055		
28							
29							
30							
31							
32							
33							
34							
35							
36							
37							
38							
39							
40							
41							
42							
43							
44							
45							
46							
47							

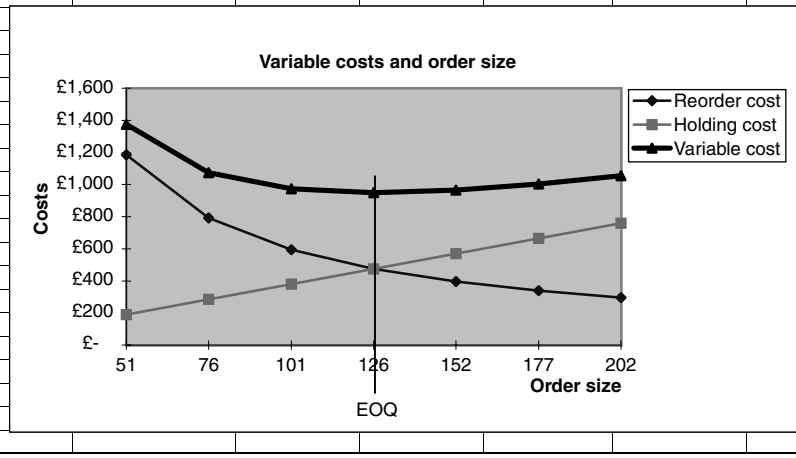


Figure 3.7 Calculating the economic order quantity

to cells D10 to D13, and the results are given in cells D15 to D18. Rows 20 to 27 show how the cost varies around the economic order quantity, and these results are then shown on a graph. Figure 3.8 gives details of the calculations.

Summary

We have built a model of an idealized inventory system that relates order size to costs and demand. This shows that large, infrequent orders have a high holding

Cell	Formula	Copied to	Explanation
D15	SQRT((2*D13*D10)/D12)		the economic order quantity
D16	D10*D11		demand * unit cost
D17	(D13*D10/D15)+(D12*D15/2)		adds reorder and holding costs
D18	D16+D17		adds all costs
B21	0.4*D15	B22:B27 (with increments of 0.2)	
C21	$D$13*$D$10/B21	C22:C27	reorder cost per period
D21	$D$12*B21/2	D22:D27	holding cost per period
E21	SUM(C21:D21)	E22:E27	total variable cost per period
F21	IF(E21=MIN($E$21:$E$27)),´minimum´)	F22:F27	identifies lowest cost in column

Figure 3.8 Calculations for the spreadsheet in Figure 3.7

cost component, so the total cost is high: small, frequent orders have a high reorder cost component, so the total cost is also high. A compromise finds the optimal order size – or economic order quantity – that minimizes inventory costs.

Review questions

3.1 What assumptions are made in calculating an economic order quantity?
3.2 What is the EOQ?
3.3 If we place orders that are larger than the EOQ, why does the total cost rise?
3.4 What does the variable cost per unit time vary with?
3.5 If we use the economic order quantity, which is bigger, the reorder cost component or the holding cost component?

Adjusting the economic order quantity

Moving away from the EOQ

Manufacturing companies often hit a problem with the EOQ. When their batch set-up costs are high, the EOQ can suggest very large batches – which complicate production scheduling, give long lead times to customers, need excessive storage capacity, and leave too much capital tied-up in stocks. These problems can be avoided by putting an artificially high value on the holding cost, but it illustrates one weakness of the calculation. Other problems arise when:

- the EOQ suggests fractional value for things which come in discrete units (an order for 2.7 lorries, for example, makes no sense and we would either buy two or three);
- suppliers are unwilling to split standard package sizes (227 kg of cement, for example, would be rounded to the nearest 50 kg);

- deliveries are made by vehicles with fixed capacities, so that 12 tonnes, say, might fit onto one lorry but the EOQ of 13 tonnes would need two lorries and hence double delivery costs;
- it is simply more convenient to round order sizes to a convenient number.

This raises the question of how much costs would rise if we do not use the EOQ. Suppose we have the following values:

- Demand, D = 6,000 units a year
- Unit cost, UC = £30 a unit
- Reorder cost, RC = £125 an order
- Holding cost, HC = £7 a unit a year

Substituting these values into the standard equations gives:

$$\text{Qo} = \sqrt{\frac{2 \times \text{RC} \times \text{D}}{\text{HC}}} = \sqrt{\frac{2 \times 125 \times 6{,}000}{7}} = 462.91 \text{ units}$$

$$\text{VCo} = \text{HC} \times \text{Q} = 7 \times 462.91 = £3{,}240.37 \text{ a year.}$$

It is unlikely that anyone would order 463 units (and certainly not 462.91), so it would be useful to know how sensitive the cost is to small changes around Qo. If we move a small distance away from the EOQ, does the cost rise very quickly, or is it relatively stable and only give small penalties? What happens here if we buy in batches of, say, 450 or 500 units?

Orders of 450 units

$$\text{VC} = \frac{\text{RC} \times \text{D}}{\text{Q}} + \frac{\text{HC} \times \text{Q}}{2} = \frac{125 \times 6{,}000}{450} + \frac{7 \times 450}{2} = £3{,}241.67$$

Orders of 500 units

$$\frac{\text{VC}}{\text{Q}} = \frac{\text{RC} \times \text{D}}{2} + \text{HC} \times \text{Q} = \frac{125 \times 6{,}000}{500} + \frac{7 \times 500}{2} = £3{,}250.00$$

Batches of 450 units – which are 2.8 per cent below optimal – raise variable costs by £1.30 or 0.04 per cent: batches of 500 units – which are 8 per cent above optimal – raise variable costs by £9.63 or 0.3 per cent. In this case the variable cost is clearly stable around the optimal value. In practice, this is always true, and we can move some way away from the EOQ and not get a significantly higher cost. We can demonstrate this by comparing the minimum variable cost, VCo, from ordering batches of size Qo, with the variable cost, VC, of ordering any other quantity, Q.

We know that:

$$\text{VCo} = \text{HC} \times \text{Qo} \quad \text{and} \quad \text{VC} = \frac{\text{RC} \times \text{D}}{\text{Q}} + \frac{\text{HC} \times \text{Q}}{2}$$

If we take the ratio of these we get:

$$\frac{VC}{VCo} = \frac{RC \times D}{Q \times HC \times Qo} + \frac{HC \times Q}{2 \times HC \times Qo}$$

Substituting $Qo = \sqrt{(2 \times RC \times D/HC)}$ gives the standard result that:

$$\boxed{\frac{VC}{VCo} = \frac{1}{2} \times \left[\frac{Qo}{Q} + \frac{Q}{Qo}\right]}$$

Now we can see the effect of using other values of Q. If we can tolerate a variable cost within, say, 5 per cent of optimal, how much will this allow us to vary the order quantity? We can find the answer by substituting $VC = 1.05 \times VCo$ in this equation and expressing Q as a fraction of Qo (so $Q = k \times Qo$). Then we have:

$$\frac{1.05 \times VCo}{VCo} = \frac{1}{2} \times \left[\frac{Qo}{k \times Qo} + \frac{k \times Qo}{Qo}\right]$$

$$1.05 = \frac{1}{2} \times \left[\frac{1}{k} + \frac{k}{1}\right]$$

$$\text{or } 2.1 \times k = 1 + k^2$$

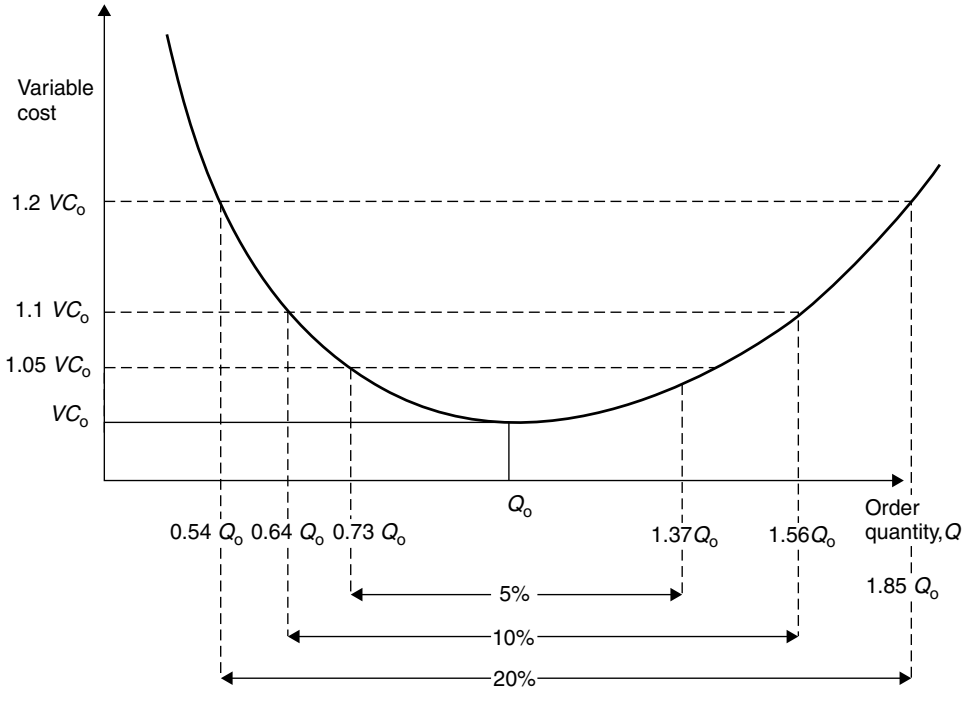


Figure 3.9 Showing the shallow variable cost curve around the EOQ

which we can solve to give two values of:

$$k = 1.37 \text{ or } k = 0.73$$

In other words, we can place orders for up to 37 per cent more than Qo, or down to 27 per cent less than Qo, and still keep the variable cost within 5 per cent of optimal (as shown in Figure 3.9). This suggests that the variable cost is very stable around the EOQ. We can get more evidence for this by accepting a variable cost within 10 per cent optimal and repeating the calculation with VC = 1.1 × VCo. Then we find the two values of k = 1.56 or 0.64, so we can place orders up to 56 per cent more than Qo, or down to 36 per cent less than Qo, and still keep variable costs within 10 per cent of optimal. This shows why the economic order quantity calculation is so useful. The analysis is based on a number of assumptions and approximations, but it gives a good guideline, and provided we use an order size that is somewhere around the EOQ, costs are likely to be close to optimal.

Worked example

Each unit of an item costs a company £40 with annual holding costs of 18 per cent of unit cost for interest charges, 1 per cent for insurance, 2 per cent allowance for obsolescence, £2 for building overheads, £1.50 for damage and loss and £4 miscellaneous costs. If the annual demand for the item is constant at 1,000 units and each order costs £100 to place, calculate the economic order quantity and the total cost of stocking the item. If the supplier will only deliver batches of 250 units, how does this affect the costs?

Solution

The annual holding cost is $18 + 1 + 2 = 21$ per cent of unit cost plus a fixed amount of £2 + £1.50 + £4 = £7.50 a year. Then:

$$D = 1{,}000 \text{ units a year}$$

$$UC = £40 \text{ a unit}$$

$$RC = £100 \text{ an order}$$

$$HC = (0.21 \times 40) + 7.50 = £15.90 \text{ a unit a year.}$$

Substituting these values gives:

$$Qo = \sqrt{\frac{2 \times RC \times D}{HC}} = \sqrt{\frac{2 \times 100 \times 1{,}000}{15.9}} = 112.15 \text{ units}$$

$$VCo = HC \times Qo = 15.9 \times 112.15 = £1{,}783 \text{ a year}$$

If Q must equal 250 units, we can find the new variable cost from:

$$\frac{VC}{VCo} = \frac{1}{2} \times \left[\frac{Qo}{Q} + \frac{Q}{Qo}\right]$$

$$VC = \frac{1,783}{2} \times \left[\frac{112.15}{250} + \frac{250}{112.15}\right] = £2,388 \text{ a year}$$

Worked example

Jessica Choi works in her bakery for 6 days a week for 49 weeks a year. Flour is delivered directly with a charge of £7.50 for each delivery. Jessica uses an average of 10 sacks of whole-grain flour a day, for which she pays £12 a sack. She has an overdraft at the bank which costs 12 per cent a year, with spillage, storage, loss and insurance costing 6.75 per cent a year.

(a) What size of delivery should Jessica use and what are the resulting costs?

(b) How much should she order if the flour has a shelf-life of 2 weeks?

(c) How much should she order if the bank imposes a maximum order value of £1,500?

(d) If the mill only delivers on Mondays, how much Jessica order and how often?

Solution

Working in consistent units we have:

$$D = 10 \times 6 \times 49 = 2,940 \text{ sacks a year}$$

$$UC = £12 \text{ a sack}$$

$$RC = £7.50 \text{ an order}$$

$$HC = (0.12+0.0675) \times 12 = £2.25 \text{ a sack a year.}$$

(a) Substituting these values:

$$Qo = \sqrt{\frac{2 \times RC \times D}{HC}} = \sqrt{\frac{2 \times 7.50 \times 2,940}{2.25}} = 140 \text{ units}$$

$$VCo = HC \times Qo = 2.25 \times 140 = £315$$

(b) If flour only has a shelf-life of 2 weeks, buying more than $2 \times 6 \times 10 = 120$ sacks at a time will mean the excess is wasted. The further the actual order size is away from EOQ, the higher the cost will be, so Jessica should aim to

have batches as near to 140 sacks as possible. In this case she should buy batches of 120 sacks. Then:

$$VC = \frac{VCo}{2} \times \left[\frac{Qo}{Q} + \frac{Q}{Qo}\right] = \frac{315}{2} \times \left[\frac{140}{120} + \frac{120}{140}\right] = £318.75 \text{ a year}$$

(c) A maximum order value of £1,500 implies 125 sacks. Again, Jessica should aim to buy as near to 140 as possible, in this case 125 sacks. Then:

$$VC = \frac{VCo}{2} \times \left[\frac{Qo}{Q} + \frac{Q}{Qo}\right] = \frac{315}{2} \times \left[\frac{140}{125} + \frac{125}{140}\right] = £317.03 \text{ a year}$$

(d) Jessica should aim for deliveries of 140 sacks. This is between 2 and 3 weeks' supply, so she has a choice of ordering every 2 weeks or ordering every 3 weeks. The cost for each of these are:

- *order every 2 weeks:* $Q = 2 \times 6 \times 10 = 120$, and we have already seen in part (b) that this gives VC = £318.75.
- *order every 3 weeks:* $Q = 3 \times 6 \times 10 = 180$, and then:

$$VC = \frac{VCo}{2} \times \left[\frac{Qo}{Q} + \frac{Q}{Qo}\right] = \frac{315}{2} \times \left[\frac{140}{180} + \frac{180}{140}\right] = £325 \text{ a year}$$

It is clearly cheaper for Jessica to order every 2 weeks.

We can again do these calculations in a spreadsheet, as shown in Figure 3.10. Here the basic data is in cells D8 to D13, while cell G9 shows the maximum percentage that we want the variable cost to be away from the optimal. Cells D15 to D28 show some results based on these figures. In particular, cell D16 shows the EOQ and D27 the corresponding variable cost. Cells G17 and I17 show the range of order size that will give variable cost within 10 per cent (in this case) of the optimal. The calculations in the spreadsheet are shown in Figure 3.11.

Orders for discrete items

One specific problem with the EOQ occurs with discrete items. If, for example, the EOQ suggests an order for 5.5 computer systems, we obviously cannot place this order and must, presumably, order either 5 or 6. This sort of problem is likely to occur when dealing with small numbers of expensive items, such as machines, engines, vehicles and drums of chemicals. We could simply round to the nearest integer, but a more formal approach asks if it is better to round up or to round down. We already know that costs rise slowly around EOQ, so this is usually not an important question. Sometimes, however, the answer makes a significant difference.

Suppose we calculate the optimal order size as Qo, which is between the integers Q′−1 and Q′ (as shown in Figure 3.12). We should round up the order size if the

	A	B	C	D	E	F	G	H	I
1	**Inventory Control-Economic Order Quantity**								
2									
3	**Notes:**	Calculations for the economic order quantity,							
4		including cost sensitivity.							
5		All units are consistent, so a lead time of one							
6		week appears as 1/52 years.							
7									
8	**Inputs**	Demand		2400		**Variable cost range**			
9		Unit cost		£ 100.00		Percentage	10%		
10		Holding cost amount		£ 10.00		Proportion	1.1		
11		Holding cost percentage		22%					
12		Reorder cost		£ 250.00					
13		Lead time		0.019230769					
14									
15	**Results**	Holding cost		£ 32.00		**Order quantity range**			
16	**Orders**	Order size		193.65		Proportion, f	0.64	to	1.56
17		Time between orders		0.08		Order size	124.27	to	301.76
18		Number of orders a period		12.39					
19		Reorder level		46.15					
20	**Investment**	Average stock		96.82					
21		Average investment		£ 9,682.46					
22		Minimum stock		0					
23		Minimum investment		£ -					
24		Maximum stock		193.65					
25		Maximum investment		£ 19,364.92					
26	**Costs**	Fixed unit costs		£ 240,000.00					
27		Variable cost		£ 6,196.77					
28		Total cost a period		£ 246,196.77					

Figure 3.10 Showing the variation in cost with order size

variable cost of ordering Q′ units is less than the variable cost of ordering Q′−1 units. That is:

$$\frac{RC \times D}{Q'} + \frac{HC \times Q'}{2} \leq \frac{RC \times D}{Q' - 1} + \frac{HC \times (Q' - 1)}{2}$$

We can simplify this by multiplying both sides by $2 \times Q' \times (Q' - 1)$.

$$HC \times Q'^2 - HC \times Q' \leq 2 \times RC \times D$$

Then we can simplify the result to:

$$Q' \times (Q' - 1) \leq Qo^2$$

This suggests a procedure for checking whether it is better to round up or round down discrete order quantities:

1. Calculate the EOQ, Qo.
2. Find the integers Q′ and Q′ − 1 that surround Qo.
3. If $Q' \times (Q' - 1)$ is less than or equal to Qo^2, order Q′.
4. If $Q' \times (Q' - 1)$ is greater than Qo^2, order Q′ − 1.

Cell	Formula	Explanation
D15	D10 + D11*D9	total holding cost
D16	SQRT((2*D12*D8)/D15)	economic order quantity
D17	D16/D8	T = Q/D
D18	1/D17	1/T
D19	D13*D8	Lead time * demand
D20	(D22 + D24)/2	average of maximum and minimum levels
D21	D20*D9	average stock * unit cost
D22	0	taken as zero for EOQ model
D24	D16	taken as Q for EOQ model
D25	D24*D9	order size, Q*unit cost
D26	D8*D9	demand * unit cost
D27	(D12*D8/D16) + (D15*D16/2)	sum of reorder and holding costs
D28	D27 + D26	total costs
G16	(($G$10*2)-SQRT(($G$10*2)^2-4))/2	first solution of quadratic equation for f
G17	G16*$D$16	first value of f*EOQ
I16	(($G$10*2) + SQRT(($G$10*2)^2-4))/2	second solution of quadratic equation for f
I17	I16*$D$16	second value of f*EOQ

Figure 3.11 Calculations for the spreadsheet in Figure 3.10

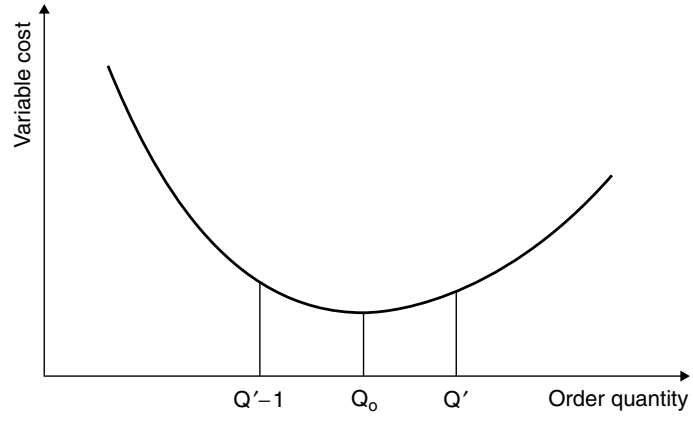


Figure 3.12 Alternative integer values around EOQ

Worked example

Schlessinger Aeronautic work a 50-week year and stock an electric motor with the following characteristics:

$$D = 20 \text{ a week}$$
$$UC = £2{,}500 \text{ a unit}$$
$$RC = £50$$
$$HC = £660 \text{ a unit a year}$$

What is the optimal order quantity? Would it make much difference if this number were rounded up or down to the nearest integer?

Solution

We can start by substituting the values to find an EOQ:

$$Qo = \sqrt{\frac{2 \times RC \times D}{HC}} = \sqrt{\frac{2 \times 50 \times 20 \times 50}{660}} = 12.31 \text{ units}$$

The company obviously cannot order 12.31 motors, so its alternatives are to order 12 or 13. Here Q′ equals 13 and the rule we developed above suggests ordering 13 when:

$$Q' \times (Q' - 1) \leq Qo^2$$

i.e.

$$13 \times 12 \leq 12.31^2$$
$$156 \leq 151.54$$

This is not true so the best policy is to order motors in batches of 12. We can check this decision by calculating the variable cost of ordering in batches of 12 or 13:

- order size 12 motors:

$$VC = \frac{RC \times D}{Q} + \frac{HC \times Q}{2} = \frac{50 \times 1{,}000}{12} + \frac{660 \times 12}{2} = £8{,}126.67$$

- order size 13 motors:

$$VC = \frac{RC \times D}{Q} + \frac{HC \times Q}{2} = \frac{50 \times 1{,}000}{13} + \frac{660 \times 13}{2} = £8{,}136.15$$

As expected, the difference is small (about 0.1 per cent).

Summary

One of the main strengths of the EOQ is that the variable cost rises slowly for orders near the optimal. The EOQ gives a good guideline, but if it cannot be used, a close approximation should give reasonable results. There is a simple procedure for checking whether it is better to round up or round down order sizes for discrete items.

Review questions

3.6 If an economic order quantity is calculated, but an order is then placed which is smaller than this, will the variable cost increase or decrease?

3.7 If we place orders that are bigger than the EOQ, the variable cost rises quickly because of the increased holding cost. Do you think that this is true?

3.8 If items have to be ordered in integer quantities, it is always better to round the EOQ upwards. Do you agree with this?

Uncertainty in demand and costs

Errors in parameters

In our derivation of the EOQ we assumed that the costs and demand were known with certainty. In practice, this is unlikely to be true. Few organizations know exactly what demand they have to meet in the future, and their costs are based on estimates and prevailing accounting conventions. As we know, the variable cost is stable around the EOQ and small errors and approximations generally make little difference. We can demonstrate this by looking at the effect on costs when there is an error in, say, forecast demand. Suppose that actual demand for an item is D, but there is a proportional error in the forecasts, E. Then the forecast is $D \times (1 + E)$ and instead of using the correct EOQ:

$$Qo = \sqrt{\frac{2 \times RC \times D}{HC}}$$

we are actually using:

$$Q = \sqrt{\frac{2 \times RC \times D \times (1 + E)}{HC}}$$

We can find the resulting error in the variable cost by substitution in:

$$\frac{VC}{VCo} = \frac{1}{2} \times \left[\frac{Qo}{Q} + \frac{Q}{Qo} \right]$$

which we can simplify to:

$$\frac{VC}{VCo} = \frac{1}{2} \times \left[\frac{1}{\sqrt{1 + E}} + \frac{\sqrt{1 + E}}{1} \right]$$

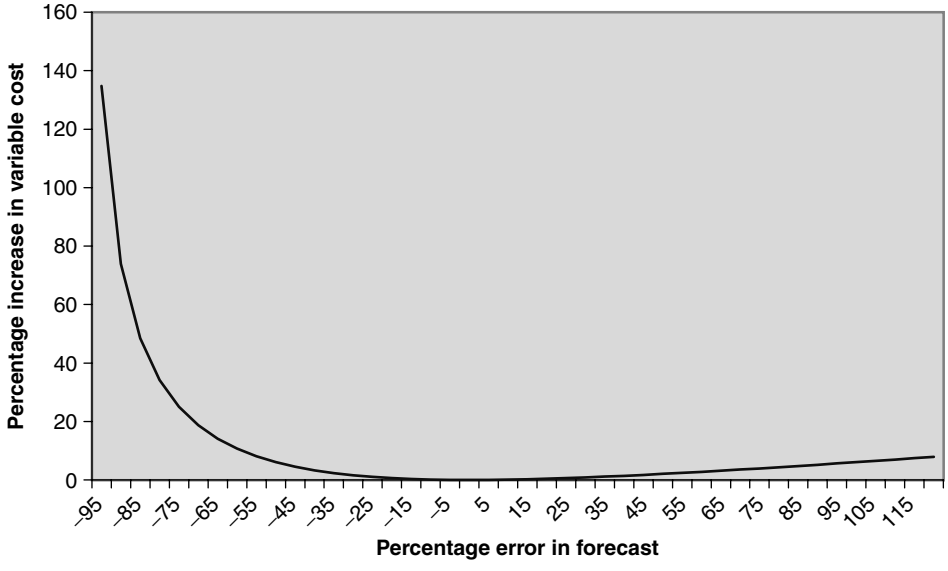


Figure 3.13 Graph of increase in variable cost against forecast error

Figure 3.13 shows the increase in variable cost against error in forecast. Again it is clear that large errors in forecast demand lead to relatively small movements away from the minimum value of VCo. Even when the forecast under-estimates demand by 50 per cent, the cost only rises by 6 per cent, and when it over-estimates demand by 50 per cent, the cost only rises by 2 per cent. This also shows that the increase in variable cost is asymmetric and rises more quickly for under-estimates of demand than for over-estimates. If, therefore, demand is uncertain and you have to make a decision with limited information, it is generally better to place larger orders than smaller ones.

We could do similar analyses to find the effect of errors in costs. Suppose, for example, that we approximate an actual reorder cost of RC by $RC \times (1 + E_1)$, and an actual holding cost of HC by $HC \times (1 + E_2)$. Then we will calculate a reorder quantity of:

$$Qo = \sqrt{\frac{2 \times RC \div (1 + E_1) \times D}{HC \times (1 + E_2)}}$$

And this gives an increase of variable cost defined by:

$$\boxed{\frac{VC}{VCo} = \frac{1}{2} \times \left[\frac{\sqrt{1 + E_1}}{1 + E_2} + \frac{\sqrt{1 + E_2}}{\sqrt{1 + E_1}} \right]}$$

Adjusting the order quantity

If costs are difficult to find, we can sometimes find values that are implied by current inventory policies. Suppose an organization is uncertain of the reorder

cost for an item. If it is already ordering the item, the actual order quantity is an estimate of the EOQ, and we can work backwards to find the implied reorder cost.

$$Qo = \sqrt{\frac{2 \times RC \times D}{HC}} \quad \text{or} \quad RC = \frac{Qo^2 \times HC}{2 \times D}$$

If this calculation is repeated for a number of items, it might be possible to get a reasonable overall estimate for the reorder cost, and this can be used to give more consistent ordering policies. We should not, however, assume that these revised calculations inevitably give better results. They can certainly be closer to a theoretical 'optimal', but inventory managers may have some reason for not accepting them, and wanting adjustments. They might, for example, know that prices are likely to rise and want to buy more than the recommended EOQ – or they may know that a marketing promotion will increase demand. There are two ways around this. First, the inventory control system could continue to present its optimal values and leave managers to adjust the actual orders. But when there is a large number of orders this will be too time-consuming. The second alternative is to have the system make these adjustments automatically. A common way of doing this is to introduce additional factors to the calculations, perhaps giving:

$$Q = K \times \sqrt{\frac{2 \times RC \times D}{HC}} \quad \text{or} \quad Q = \sqrt{\frac{2 \times RC \times D}{HC \times K}}$$

A more straightforward adjustment allows managers to use surrogate costs, which are not actual costs, but which give the type of orders that they want. If, for example, they prefer orders that are generally smaller than the EOQ, they can substitute artificially low reorder costs or high holding costs. This kind of adjustment should only be made if there are genuine reasons and after careful analysis.

Worked example

A company has a standing order of 40 units of an item every month. What can you infer about the costs? If the reorder cost is actually €160, what is the implied holding cost?

Solution

We know that the company orders 40 units every month, so we can assume that the demand is constant at around 10 units a week. If the company believes that this is the optimal ordering policy:

$$Qo = \sqrt{\frac{2 \times RC \times D}{HC}} \quad \text{or} \quad 40 = \sqrt{\frac{2 \times RC \times 10}{HC}}$$

So:

$$RC/HC = 80$$

The company believes that the reorder cost is 80 times the weekly holding cost. If the reorder cost is €160, the implied holding cost is €2 a unit a week.

Summary

We can estimate the effect of errors in the parameters used to calculate the economic order quantity. These effects are usually small. Often managers will adjust the calculated results to get values that are not 'optimal' but which give features that they want.

Review questions

3.9 Do the assumptions and other limitations of the economic order quantity calculation mean it has no practical value?

3.10 If you are uncertain about the values used to find an EOQ, would you generally prefer larger or smaller orders?

Adding a finite lead time

Causes of lead time

So far we have been looking at the question of how much to order, and assumed that lead time is zero. Then, as soon as we place an order, the materials arrive and are ready for use. In practice, this almost never happens and there can be significant delays. These delays form a lead time which is the total time between placing an order and having the materials available for use. The lead time occurs because of:

- *time for order preparation*. When an organization decides to buy something, there can be a delay before the order is ready to send to a supplier. With small orders this can be very short and involve a few administrative details, or the procedures can be completely automated. Larger orders need time for designing items, going through some kind of tendering process, arranging finance, and so on.
- *time to get the order to the right place in suppliers*. With automated ordering and e-business this is largely eliminated, but there can be delays with more traditional approaches, or when, say, the recipients are away from their offices.
- *time at the supplier*. This is the time needed for a supplier to process the order and prepare the delivery. It is highly variable and can be very short for small items already in stock, or very long for items that have to be specially designed and made.

- *time to get materials delivered from suppliers*. This can be a few hours for local suppliers, or several weeks for an international shipment, or months for a special and complex delivery.
- *time to process the delivery* is the total time between receiving a delivery and getting the materials available for use. This might include time for checking, inspection, recording, cataloguing, movement, and so on.

This lead time can vary between a few minutes to several years, and is typically between a few days and a few weeks. It is in everybody's interests to make it as short as possible – customers want their deliveries as fast as possible, and suppliers want to maintain high customer service and not keep materials in their own stocks. A lot of work in the past few years has looked for ways of reducing the lead time. With e-business, for example, most of the administration of regular orders is removed. This lowers the reorder cost and gives smaller, more frequent deliveries from suppliers. These, in turn, reduce the variability and uncertainty that are the main reasons for holding stock, and so reduce overall costs. A similar effect comes from, say, agile operations which respond quickly to customer demands, or improved transport which reduces journey times.

By streamlining their operations and administration, organizations can reduce the lead time to little more than the time needed for physical delivery. This can still be quite long, especially for international transfers, but local suppliers can give fast delivery – which is why you see car assembly plants and other major industries surrounded by their suppliers. Of course, there are many circumstances in which the lead time cannot be reduced. If, for example, you want a specialized reaction vessel for an oil refinery, you will need several years for design, negotiations, finance and special delivery. If you want some 'Jersey Royal' potatoes, you will have to wait until they are planted, grow, are harvested, sorted and delivered.

Reorder level

If we assume the lead time is constant, we can add a useful extension to our previous analysis. Consider the effect of a finite lead time, LT, on the standard stock level pattern. When demand is constant, there is no benefit in carrying stock from one cycle to the next, so each order should be timed to arrive just as existing stock runs out. To achieve this, we have to place an order a time LT before the delivery is needed (as shown in Figure 3.14). The easiest way of arranging this is to define a *reorder level*. When stock declines to this reorder level, it is time to place an order. The EOQ does not depend on lead time and remains unchanged.

We can calculate a reorder level from the following argument. When it is time to place an order, the stock on hand must just cover the demand until this order arrives. As both demand and lead time are constant, the amount of stock needed to cover the lead time is also constant at:

$$\text{lead time} \times \text{demand per unit time}$$

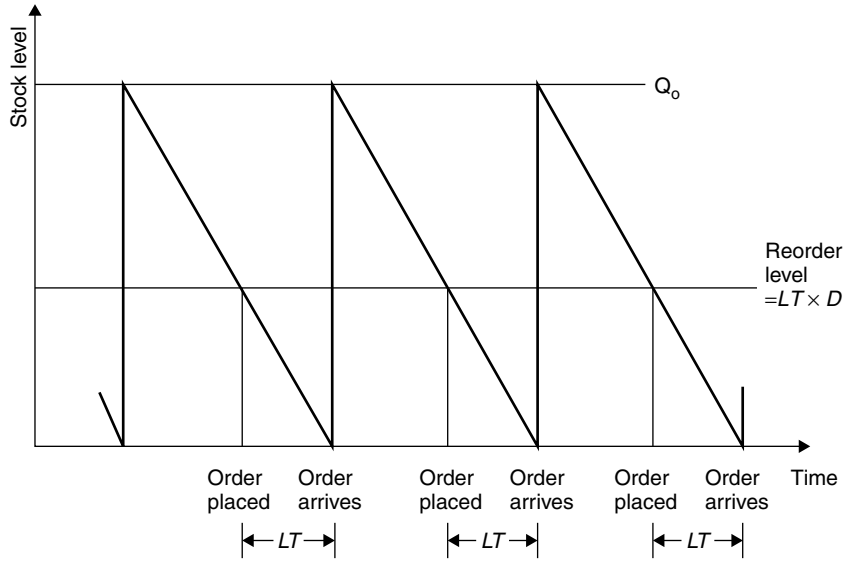


Figure 3.14 The lead time and corresponding reorder level

So this defines the reorder level:

reorder level = lead time demand = lead time × demand per unit time

$$ROL = LT \times D$$

The simple rule, then, is to order a batch of size Qo whenever the stock level falls to LT × D.

Worked example

Carl Smith uses radiators at the rate of 100 a week, and he has calculated an EOQ of 250 units. What is his best ordering policy if lead time is: (a) one week? or (b) two weeks?

Solution

(a) Substituting values LT = 1 and D = 100 gives:

$$ROL = LT \times D = 1 \times 100 = 100 \text{ units}$$

As soon as the stock level declines to 100 units, Carl should place an order for 250 units (as shown in Figure 3.15).

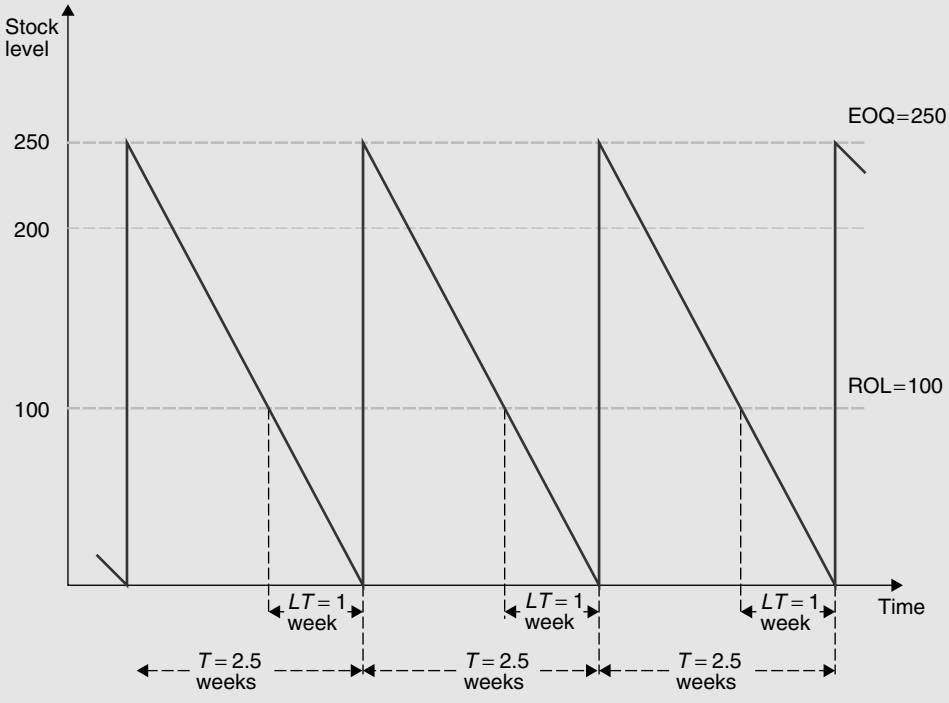


Figure 3.15 Reorder level for radiators

(b) Substituting values LT = 2 and D = 100 gives:

$$\text{ROL} = \text{LT} \times \text{D} = 2 \times 100 = 200 \text{ units}$$

As soon as the stock level declines to 200 units, Carl should place an order for 250 units.

Longer lead times

This rule of ordering when stock level declines to lead time demand works well provided the lead time is shorter than the stock cycle. We can demonstrate this with the last worked example. Here the cycle length was:

$$\text{T} = \text{Q/D} = 250/100 = 2.5 \text{ weeks}$$

Our rule worked well when the lead time was 1 and 2 weeks, but what happens when the lead time is increased to 3 weeks? Then the reorder level is:

$$\text{ROL} = \text{LT} \times \text{D} = 3 \times 100 = 300$$

But Figure 3.15 shows that the stock level varies between zero and 250 units and never actually rises to 300. The problem is that when the lead time is longer than the cycle length, there is always one order outstanding (as shown in Figure 3.16). So when it is time to place order B, there is one order, A, of 250 units outstanding and due to arrive before B. This suggests that the stock on hand plus the outstanding order must together be enough to last until B arrives – so they must add to equal the lead time demand.

$$\text{stock on hand} + \text{stock on order} = \text{LT} \times \text{D} = 300 \text{ units}$$

As there are 250 units of stock on order, we place a new order when the stock on hand falls to (300 – 250 =) 50 units. The general rule, then, is to place an order whenever the stock on hand falls to lead time demand minus the stock on order, or:

$$\text{reorder level} = \text{lead time demand} - \text{stock on order}$$

When the lead time is particularly long, there can be several orders outstanding at any time. In particular, when the lead time is between n and $n+1$ cycle lengths, giving:

$$n \times T < LT < (n+1) \times T$$

there are n orders outstanding when it is time to place another. Then we subtract $n \times Qo$ from the lead time demand to get the reorder level.

$$\text{reorder level} = \text{lead time demand} - \text{stock on order}$$
$$\text{ROL} = \text{LT} \times \text{D} - n \times Qo$$

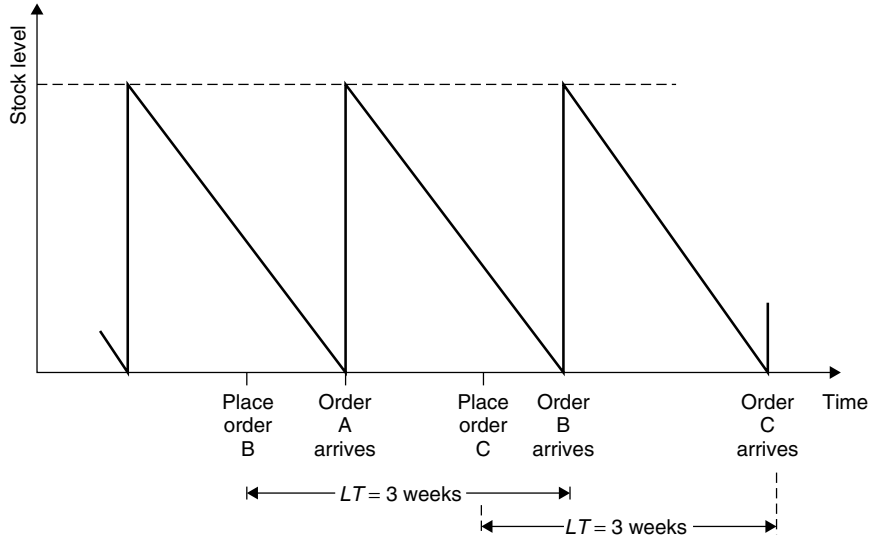


Figure 3.16 Timing when the lead time is longer than the cycle length

Worked example

Demand for an item is steady at 1,200 units a year with an ordering cost of £16 and holding cost of £0.24 a unit a year. Describe an appropriate ordering policy if the lead time is constant at (a) 3 months; (b) 9 months; or (c) 18 months.

Solution

The values given are:

$$D = 1{,}200 \text{ units a year}$$

$$RC = £16 \text{ an order}$$

$$HC = £0.24 \text{ a unit a year}$$

We can substitute these to find the optimal order size and corresponding cycle length:

$$Qo = \sqrt{\frac{2 \times RC \times D}{HC}} = \sqrt{\frac{2 \times 16 \times 1{,}200}{0.24}} = 400 \text{ units}$$

and

$$To = Qo/D = 400/1200 = 0.33 \text{ years or } 4 \text{ months}$$

With lead time between $n \times To$ and $(n + 1) \times To$ there are n orders outstanding when it is time to place the next order, so the reorder level is $LT \times D - n \times Qo$

(a) A lead time of 3 months is less than one cycle length, so $n = 0$ and the reorder level, working in months, is:

$$ROL = LT \times D = 3 \times 100 = 300 \text{ units}$$

Every time the stock on hand falls to 300 units, it is time to place an order for 400 units.

(b) A lead time of 9 months is between 2 and 3 stock cycles, so $n = 2$ and the reorder level is:

$$ROL = LT \times D - n \times Qo = 9 \times 100 - 2 \times 400 = 100 \text{ units}$$

Every time the stock on hand falls to 100 units, it is time to place an order for 400 units.

(c) A lead time of 18 months is between 4 and 5 stock cycles, so $n = 4$ and the reorder level is:

$$ROL = LT \times D - n \times Qo = 18 \times 100 - 4 + 400 = 200 \text{ units}$$

Every time the stock on hand falls to 200 units, it is time to place an order for 400 units

Some practical points

The analysis in this chapter has shown when to place an order – set by the reorder level – and how much to order – set by the economic order quantity. This gives the basis of a useful stock control system. In practice, inventories of any size are controlled by computer, but some simple procedures can be used for any store, even very small ones. A *two-bin system*, for example, gives a simple procedure for controlling stock without computers or continuous monitoring of stock levels. Stock of an item is kept in two separate bins, A and B. Bin B contains an amount equal to the reorder level, and all remaining stock is in bin A. Stock is used from bin A until it is empty – at which point only the reorder level remains and it is time to place another order. Then stock is used from bin B until the order arrives, when B is again filled with the reorder level and all extra units are put into A. This system can work well for cheap goods with long cycle times and can be modified so that the 'bins' are actually levels of liquid in containers (such as oil tanks), spaces on shelves, powders or grains in silos, and so on.

The two-bin system can be extended to a three-bin system which allows for some uncertainty. In this, the third bin holds a reserve that is only used in an emergency. Then the normal stock is used from bin A. When this is empty the reorder level in bin B shows that it is time to place another order. When bin B is empty it shows that the delivery is overdue and an urgent delivery is needed. Stock from bin C can be used until this arrives.

There can also be some practical difficulties with the reorder level. Our calculations assume that the lead time is known exactly and constant. In practice, there can be quite wide variation, allowing for availability, supplier reliability, checks on deliveries, transport conditions, customs clearance, delays in administration, and so on. The lead time may actually be difficult to find when, for example, it seems to differ from a supplier's quoted figure. If a supplier says that deliveries will be made 'within an average of two working days', this clearly depends on their definition of 'working days' and 'average'. If your delivery arrives after seven days, you might find that two of these days are the weekend, one day is a local holiday, and the remaining four days balance the four other customers who had deliveries after one day. We will look at problems with variable lead time in Chapter 5.

There can also be a problem with identifying the exact point when stock on hand actually falls to the reorder level. Some stock levels are not recorded continuously but are checked periodically, perhaps at the end of the week. Then an unexpectedly large demand might reduce stocks well below the reorder level before they are checked. Even if the stocks are checked continuously, some large orders may arrive at the same time, so that stocks instantly fall from above the reorder level to considerably below it. Another problem appears with large stocks, such as chemical tanks, coal tips or heaps of raw materials, where the stock level is only known approximately. A power station might have a stock of 6,000 tonnes of coal, but this is an approximation and nobody is actually going to weigh it. Then the reorder level might be passed without anyone noticing.

Another point is that we have assumed that the order size is independent of the lead time and the reorder level. In practice, people often prefer larger orders if there are longer lead times, and they raise the reorder level to add an element of safety. The reorder level can also influence the order size, with people typically placing smaller orders with higher reorder levels (perhaps because of a reluctance to place a large order when there are already fairly high stocks). Such decisions should, of course, be made after some careful analysis.

Summary

There is usually a lead time between placing an order and having the materials available for use. The trend is towards shorter lead times. We need some way of finding the time to place an order, and one approach uses a reorder level. Then we order an amount equal to the economic order quantity, whenever stock falls to the reorder level. For constant lead time and demand, the reorder level equals lead time demand minus any stock on order.

Review questions

3.11 What is the reorder level?

3.12 If demand for an item is 10 units a week, the economic order quantity is 30 units and the lead time is 7 weeks, how many orders will be outstanding when it is time to place another order?

3.13 Orders should be placed when stock on hand declines to the reorder level; does this mean that continuous monitoring of stock levels is needed?

Chapter review

- This chapter introduced the idea of using quantitative models for controlling stocks for independent demand. In particular, it showed how to analyse a simple inventory system and find an ordering policy that minimizes total costs.
- Large, infrequent orders have a high holding cost component, so the total cost is high: small, frequent orders have a high reorder cost component, so the total cost is also high. A compromise finds the optimal order size – or economic order quantity – that minimizes inventory costs.
- The variable cost of stock rises slowly for orders near the EOQ. The EOQ gives a good guideline, but if it cannot be used a close approximation should give reasonable results.
- There is a simple procedure for checking whether it is better to round up or round down order sizes for discrete items.
- We can estimate the effect of errors and approximations in the parameters used to calculate the EOQ. These effects are usually small.
- Often managers will adjust the calculated results to get values that are not 'optimal' but which give features that they want.

- There is usually a lead time between placing an order and having the materials available for use. The trend is towards shorter lead times.
- We can use the lead time to calculate a reorder level which shows when it is time to place an order. In particular, we order an amount equal to the economic order quantity whenever the stock level falls to the reorder level. For constant lead time and demand, the reorder level equals lead time demand minus any stock on order.

Project

There is a lot of software for doing calculations for inventory control. Most of this is found in large and comprehensive packages. We can, however, demonstrate the calculations using simple programs or even spreadsheets. Figures 3.7 and 3.10, for example, show some of the calculations for the economic order quantity on spreadsheets.

The aim of this project is to develop a spreadsheet for exploring variations in inventory costs. Have a look at the spreadsheets given, check them and the answers they give. What other features would be useful? What other calculations can you add? Develop your own spreadsheets and use them to see how costs vary for a range of changing circumstances and assumptions.

Problems

3.1 Demand for an item is constant at 1,000 units a year. Unit cost is £50, reorder cost is £100, holding cost is 25 per cent of value a year and no shortages are allowed. Describe an optimal inventory policy for the item. What order size will give a variable cost within 10 per cent of optimal? What is the cost if suppliers only make deliveries of 200 units?

3.2 A company is introducing a new item and it has forecast likely demand next year as between 100 and 130 units. The costs are uncertain, but the reorder cost is somewhere between $50 and $70, and the holding cost is between 20 per cent and 25 per cent of unit cost a year. If the unit cost is $200, what can you say about the order size?

3.3 Per Norstrom supplies computer systems to a warehouse in Rotterdam. He sells 16 systems a week. The cost of an average system is €5,000, while order administration costs and delivery from Malaysia cost €1,000. The lead time is around 4 weeks and holding costs are around 16 per cent a year. What policy would you recommend for Per?

3.4 Demand for an item is constant at 40 units a week, and the economic order quantity is calculated to be 100 units. What is the reorder level if lead time is constant at 4 weeks? What is the effect of adding some margin of safety and raising the reorder level by ten units? What happens if the lead time (a) falls to 2 weeks or (b) rises to 6 weeks?

3.5 (a) Arcadia Windings is concerned about its stocks of copper cable. The demand for this is 8,000 metres a week, with a cost of £4 a metre. Each order costs £350 for administration and £550 for delivery, and has a lead time of 8 weeks. Holding costs are about 25 per cent of value held a year, and any shortages would disrupt production and give very high costs. What is the best inventory policy for the cable? How does this compare with the current policy of placing a regular order every week?

(b) How would this analysis differ if the company wanted to maximize profit rather than minimize costs? What is the gross profit if the company sells cable for £8 a metre?

Discussion questions

3.1 The main concern of an organization is customer service, and this often needs high stocks. Why, then, are we concentrating on minimizing the cost of stock, when we should be more concerned with raising stock levels?

3.2 What costs are incurred by holding stock? How would you set about finding these? Why are shortage costs so difficult to find?

3.3 What are the assumptions of the economic order quantity? How valid are these? What factors in real inventory control are not included in the economic order quantity model?

3.4 The variable cost rises slowly around the EOQ, and the analysis is based on a series of assumptions and approximations. Why, then, do we bother with the calculations rather than allowing inventory managers to design policies based on their experience? Would they get good results without bothering with the formal analysis?

3.5 What are the benefits of short lead times? How can these be achieved in practice?

References and further reading

Greene, J.H. (1997) *Production and Inventory Control Handbook* (3rd edition), New York: McGraw-Hill.

Harris, F. (1915) *Operations and Cost*, Chicago: A.W. Shaw Co Factory Management Series.

Institute of Management and Administration (2002) *The IOMA Handbook*, Chichester: John Wiley.

Lewis, C.D. (1997) *Demand Forecasting and Inventory Control*, Cambridge: Woodhead Publishing.

Muller, M. (2002) *Essentials of Inventory Management*, London: Amacom.

Tersine, R.J. (1994) *Principles of Inventory and Materials Management*, Englewood Cliffs, NJ: Prentice Hall.

Waters, D. (2002) *Operations Management Producing Goods and Services*, Harlow: Pearson Education.

Wilson, R.H. (1934), A Scientific Routine for Stock Control, Harvard Business Review XIII.

Zipkin, P. (2000) *Foundations of Inventory Management*, Homewood, IL: Irwin.

4

Models for Known Demand

Aims of the chapter

The last chapter described the economic order quantity. This analysis is based on a number of assumptions. In the next two chapters we describe some models where these assumptions are removed. This chapter keeps the condition that all variables take known values (giving deterministic models) while the next chapter allows some uncertainty (giving probabilistic models).

Hundreds of models have been developed for different circumstances, so it is impossible to describe even a fraction of these. In these two chapters we concentrate on the most widely used models, and use these to illustrate some general principles.

After reading this chapter you should be able to do the following:

- calculate the best order size when there are discounts in the unit cost;
- extend this analysis to problems that have discrete changes in other costs;
- calculate optimal batch sizes when there is a finite replenishment rate;
- appreciate the use of planned shortages;
- find optimal policies for back-orders;
- design a policy for maximizing revenue with lost sales;
- solve problems with constraints on space or investment;
- calculate ordering policies with discrete, variable demand.

This chapter emphasizes:

- variables that are known with *certainty*;
- *deterministic models* for inventory control;
- *removing assumptions* of the economic order quantity.

Price discounts from suppliers

Variable costs

In the last chapter we assumed that all costs are fixed – so they have constant, known values that never change. In this chapter we start by seeing what happens when the costs vary with the quantity ordered. You can often see this with discounted unit prices, where a supplier quotes lower prices for larger orders. A particular computer, for example, might cost £2,500, but this falls to £2,250 for orders of ten or more, and to £2,000 for orders of 50 or more.

The reorder cost might also vary with order size, particularly if it includes quality control inspections, or some other processing of deliveries. The transport cost also varies with the amount delivered, and can have quite complex patterns. For example, orders up to a certain size might fit into a standard delivery vehicle or container, but larger orders need a second vehicle which doubles the delivery cost. Or orders above a certain size might switch from airfreight to sea, reducing the unit cost but increasing the lead time.

Even the holding cost can vary with order size. If, for example, the economic order quantity is large, there may not be enough warehouse space to receive a full order and extra space must be rented. Overall, then, it is common for all of the costs to vary, often in quite complicated ways. We could build models to investigate the patterns of all of these variations, but they would follow similar paths. We will, therefore, illustrate the general approach with problems where the unit cost decreases in discrete steps with increasing order quantity.

Valid total cost curve

The most common variation in cost occurs when a supplier offers a reduced price on all units for orders above a certain size. There is often more than one discounted price, giving the pattern of unit cost shown in Figure 4.1. The basic unit cost is UC_1, but this reduces to UC_2 for orders bigger than Q_a, to UC_3, for orders bigger than Q_b, to UC_3 for orders bigger than Q_c, and so on.

unit cost	Order quantity lower limit	upper limit
UC_1	0	Q_a
UC_2	Q_a	Q_b
UC_3	Q_b	Q_c
UC_4	Q_c	Q_d
UC_5	Q_d	Q_e
etc.		

If we look at the most expensive unit cost, UC_1, we can draw a graph of the total cost per unit time against the order size, as we did to find the economic order quantity. In this case, though, the curve will only be valid for order quantities in

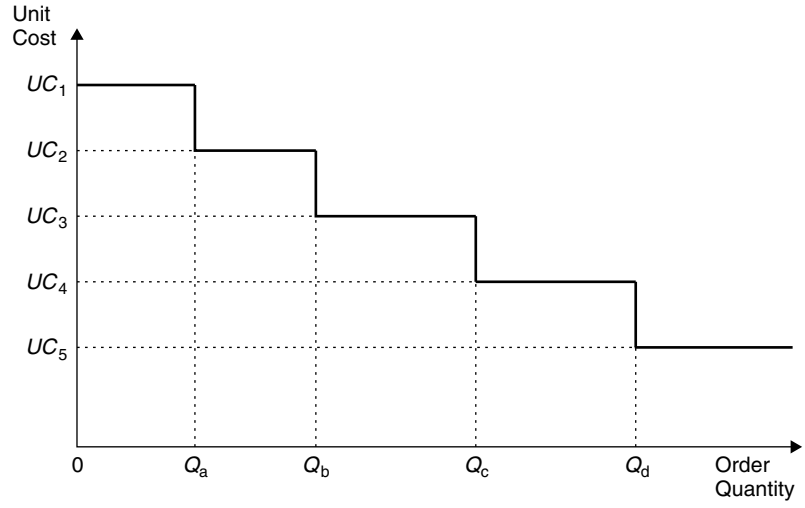


Figure 4.1 Stepped decrease in unit cost with reorder quantity

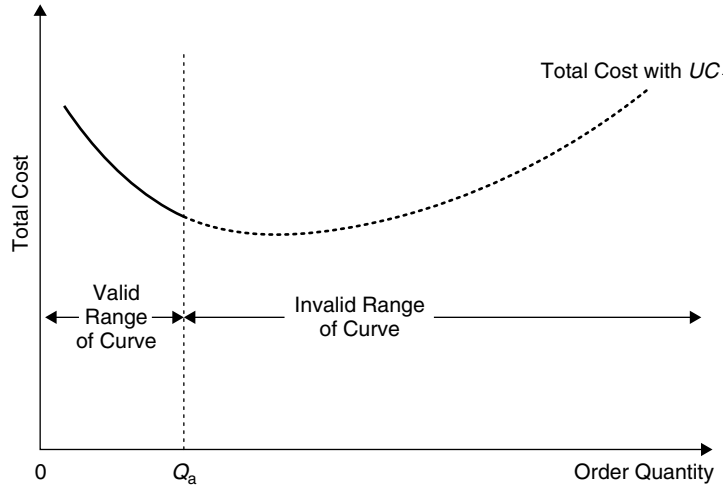


Figure 4.2 The valid range of the total cost curve using UC_1

the range zero to Q_a. In Figure 4.2 the continuous line shows the section of the total cost curve that applies to this unit cost, while the broken line shows the part of the curve that does not apply. If we place an order for some quantity that is not between 0 and Q_a the unit cost will be different and we should be working on a different total cost curve.

If we now consider the next highest unit cost, UC_2, we can draw a second total cost curve. As UC_2 is less than UC_1, this second curve will always be lower than the original one. It will only be valid in the range Q_a to Q_b, and if we place orders

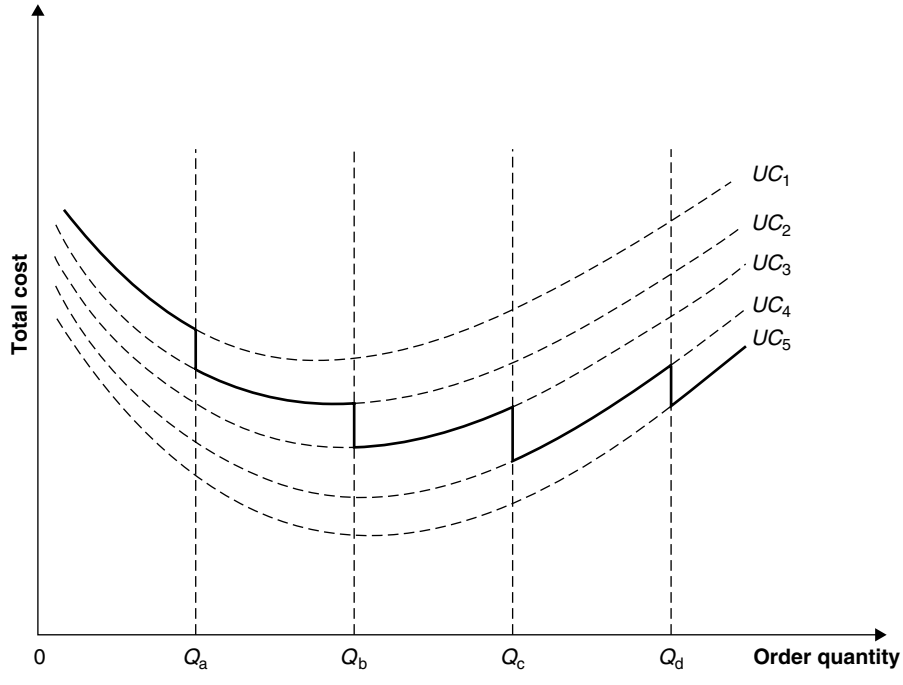


Figure 4.3 A valid cost curve with five unit costs

with a size outside this range we should be working on a different curve. Now, we can take the next highest unit cost UC_3 and draw a third total cost curve below the first two and valid in the range Q_b to Q_c. Eventually, we have a total cost curve for each unit cost, as shown in Figure 4.3. These curves never cross each other but remain separate, and each is valid within a prescribed range. If we join up the valid sections, we get an overall graph of the total cost per unit time, shown by the continuous line of the *valid cost curve*.

Finding the lowest valid cost

We want to find the order quantity that minimizes the total cost per unit time. In other words, we are looking for the optimal value of Q that corresponds to the lowest point on the valid cost curve.

We know from Chapter 3 that in general:

$$\text{Qo} = \sqrt{\frac{2 \times \text{RC} \times \text{D}}{\text{HC}}}$$

We can express the holding cost as a proportion, I, of the unit cost, and for each unit cost UC_i, the minimum point of the cost curve comes with Qo_i. Then Qo_1 gives the lowest point on the total cost curve for UC_1, Qo_2 gives the lowest point

on the total cost curve for UC_2, and so on. We also know that:

$$Qo_i = \sqrt{\frac{2 \times RC \times D}{I \times UC_i}}$$

For each curve with unit cost UC_i this minimum is either 'valid' or 'invalid':

- *A valid minimum* is within the range of valid order quantities for this particular unit cost.
- *An invalid minimum* falls outside the valid order range for this particular unit cost.

Every set of cost curves will have at least one valid minimum, and a variable number of invalid minima, as shown in Figure 4.4.

You can see two other interesting features in the valid cost curve. First, the valid total cost curve always rises to the left of a valid minimum. This means that when we search for an overall minimum cost it is either *at* the valid minimum or somewhere to the right of it. Second, there are only two possible positions for the

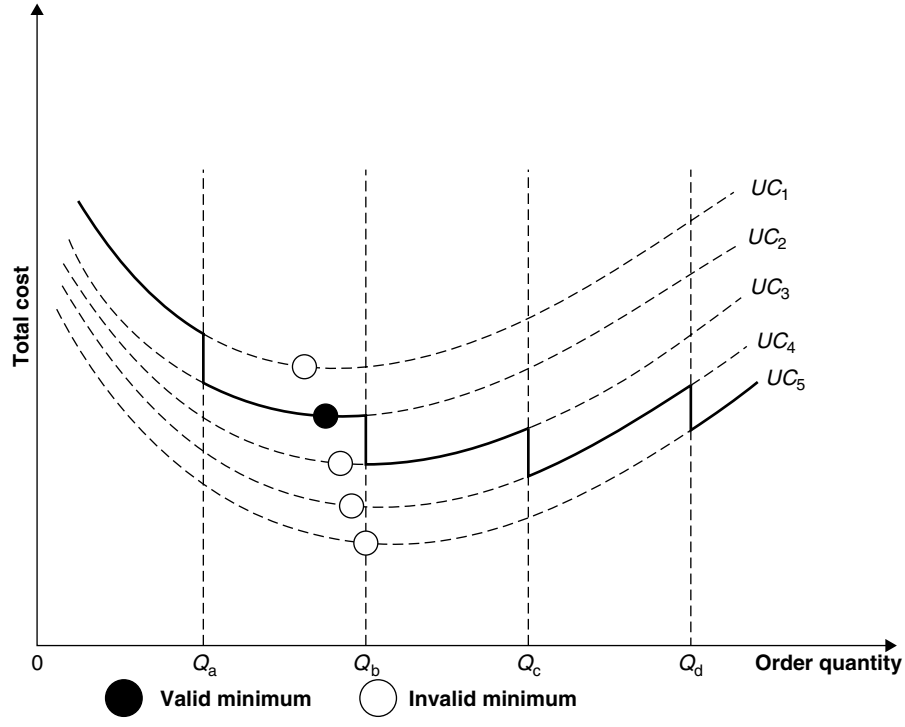


Figure 4.4 Valid and invalid minima for cost curves

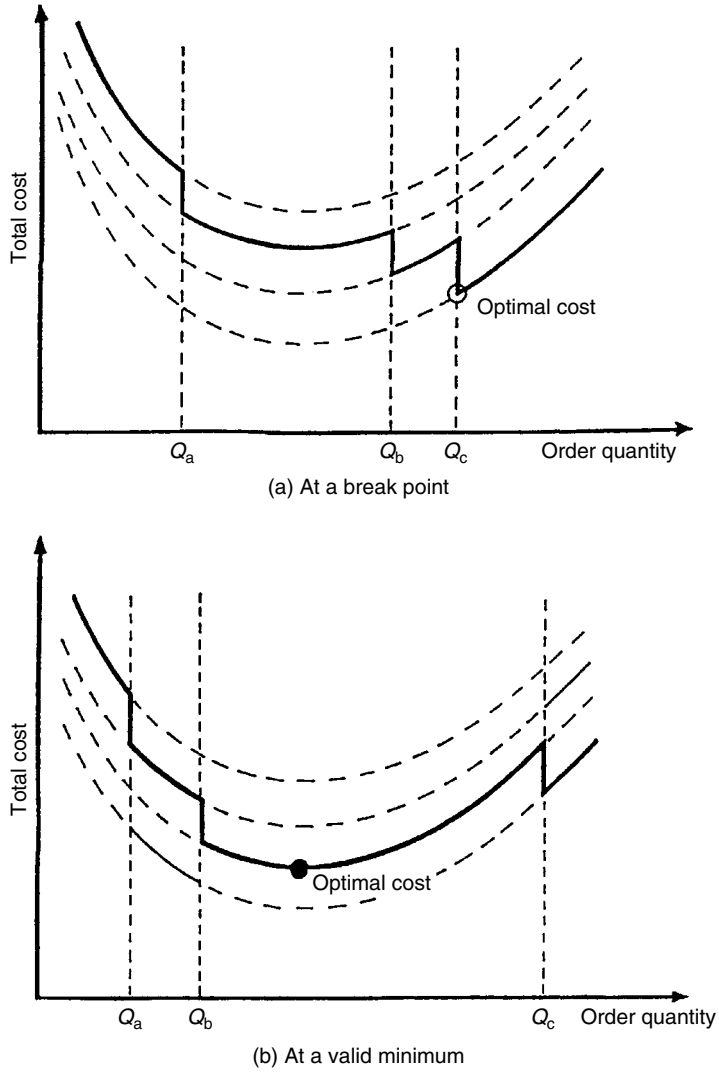


Figure 4.5 Possible positions of the lowest cost

overall minimum cost: it is either at a valid minimum, or else at a cost break point (as shown in Figure 4.5).

These observations suggest an efficient way of finding the lowest point on the valid cost curve. All we have to do is find the cost at the valid minimum and compare this with the cost at each break point to the right of this valid minimum. The optimal order quantity corresponds to the lowest of these costs. A formal description of this procedure is shown in the flow chart of Figure 4.6. Starting with the lowest unit cost we find the minimum point on the corresponding total cost curve. If this is valid, it must be the optimal solution: if it is invalid, we find the

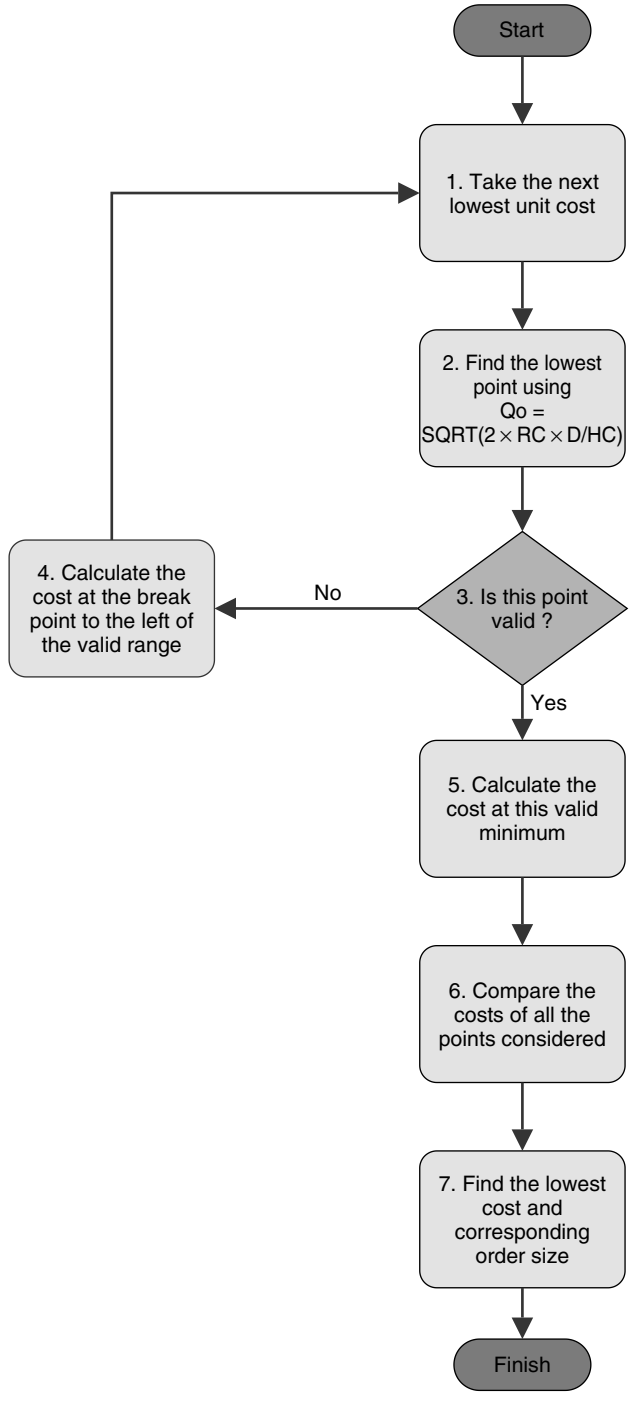

Start
1. Take the next lowest unit cost
2. Find the lowest point using Qo = SQRT(2 × RC × D/HC)
3. Is this point valid ?
No
4. Calculate the cost at the break point to the left of the valid range
Yes
5. Calculate the cost at this valid minimum
6. Compare the costs of all the points considered
7. Find the lowest cost and corresponding order size
Finish

Figure 4.6 Procedure for finding the best order size

total cost at the break point at the smaller end of the valid range. We repeat this procedure for each cost curve in turn until we find a valid minimum. Then we have a set of costs – one from each unit cost curve – and we can compare these and identify the lowest, and the corresponding order size.

Worked example

The annual demand for an item is 2,000 units, each order costs £10 and the annual holding cost is 40 per cent of unit cost. The unit cost depends on the quantity ordered as follows:

- £1 for order quantities less than 500
- £0.80 for quantities between 500 and 999
- £0.60 for quantities of 1,000 or more.

What is the optimal order size?

Solution

Listing the variables we know:

$$D = 2{,}000 \text{ units a year}$$
$$RC = £10 \text{ an order}$$
$$I = 40 \text{ per cent of unit cost a year}$$
$$HC = I \times UC = 0.4 \times UC \text{ a unit a year}$$
$$UC \text{ varies and is either } £1, £0.80 \text{ or } £0.60$$

We should also remember that the total cost, TCo, for an optimal order quantity, Qo, is:

$$TCo = UC \times D + HC \times Qo$$

While the total cost, TC, for any other order quantity, Q, is:

$$TC = UC \times D + \frac{RC \times D}{Q} + \frac{HC \times Q}{2}$$

Now we can follow the procedure given in Figure 4.6.
Taking the lowest cost curve:

- UC = £0.60, valid for Q of 1,000 or more
- $Qo = \sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 10 \times 2{,}000/0.4 \times 0.60)} = 408.2$

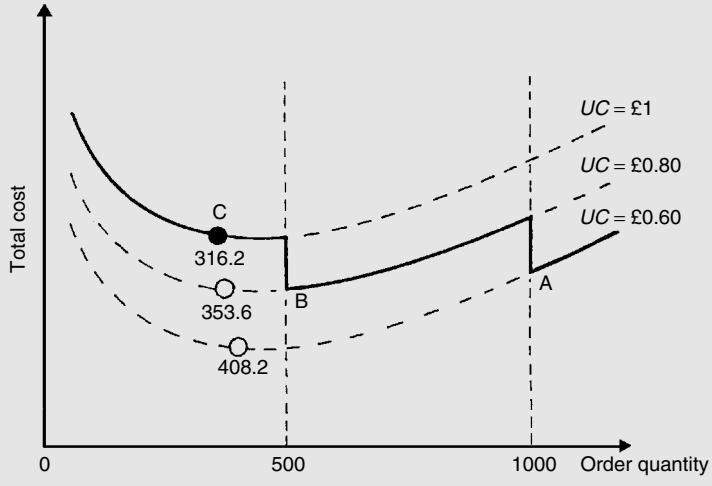


Figure 4.7 Valid total cost curve for the worked example

- This is an *invalid* minimum as Qo is not greater than 1,000.
- Calculate the cost of the break-point at the lower end of the valid range (i.e. Q = 1,000):

$$\begin{aligned} TC &= UC \times D + (RC \times D)/Q + (HC \times Q)/2 \\ &= 0.60 \times 2{,}000 + (10 \times 2{,}000)/1{,}000 + (0.4 \times 0.60 \times 1{,}000)/2 \\ &= £1{,}340 \text{ a year (point A in Figure 4.7).} \end{aligned}$$

Taking the next lowest cost curve:

- UC = £0.80, valid for Q between 500 and 1,000
- $Qo = \sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 10 \times 2{,}000/0.4 \times 0.80)} = 353.6$
- This is an *invalid* minimum as Qo is not between 500 and 1,000.
- Calculate the cost of the break-point at the lower end of the valid range (i.e. Q = 500):

$$\begin{aligned} TC &= UC \times D + (RC \times D)/Q + (HC \times Q)/2 \\ &= 0.80 \times 2{,}000 + (10 \times 2{,}000)/500 + (0.4 \times 0.80 \times 500)/2 \\ &= £1{,}720 \text{ a year (point B in Figure 4.7).} \end{aligned}$$

Taking the next lowest cost curve:

- UC = £1.00 valid for Q less than 500
- $Qo = \sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 10 \times 2{,}000/0.4 \times 1.00)} = 316.2$

- This is a *valid* minimum as Qo is less than 500.
- Calculate the cost at this valid minimum:

$$\begin{aligned} TCo &= UC \times D + HC \times Qo \\ &= 1 \times 2{,}000 + 0.4 \times 1.00 \times 316.2 \\ &= £2{,}126.48 \text{ a year(point C in Figure 4.7)} \end{aligned}$$

Even if there were more cost curves we need continue no further. Once we have found a valid minimum (point C in this case), we can stop looking and compare the solutions obtained so far. The lowest cost corresponds to the best order size. The choice here is between:

point A	Q = 1,000	cost = £1,340 p.a.
point B	Q = 500	cost = £1,720 p.a.
point C	Q = 316.2	cost = £2,126.49 p.a.

The best choice is clearly to order batches of 1,000 units, placing one order every six months and with total annual costs of £1,340. The discounts offered on unit cost more than offset the increased cost of carrying extra stock.

Worked example

During a 50-week year IH James and Partners notice that demand for one of its products is more or less constant at 10 units a week. The cost of placing an order, including delivery, is around €150. The company aims for 20 per cent annual return on assets. The supplier of the item quotes a basic price of €250 a unit, with discounts of 10 per cent on orders of 50 units or more, 15 per cent per cent on orders of 150 units or more and 20 per cent on orders of 500 units or more. What is the optimal order size for the item?

Solution

We can list the variables as:

$$D = 10 \times 50 = 500 \text{ units a year}$$

$$RC = €150 \text{ an order}$$

$$I = 20 \text{ per cent of unit cost a year}$$

$$HC = I \times UC = 0.2 \times UC \text{ a unit a year}$$

We can do the calculations for this in a spreadsheet, as illustrated in Figure 4.8. This shows the total cost associated with each unit cost, and identifies the best policy as ordering 150 units every 15 weeks. The valid cost curve is shown in Figure 4.9.

	A	B	C	D	E	F
1	**Inventory Control - Economic Order Quantity with discounts**					
2						
3	Note:	Adding discounts in unit price				
4		to the economic order quantity				
5						
6	**Inputs**	Demand		500		
7		Holding cost percentage		20%		
8		Reorder cost		€ 150.00		
9						
10		Unit cost	250.0	225.0	212.5	200
11		For orders over	0	50	150	500
12						
13	**Results**	Unit cost	€ 250.00	€ 225.00	€ 212.50	€ 200.00
14		Holding cost	€ 50.00	€ 45.00	€ 42.50	€ 40.00
15		EOQ	54.8	57.7	59.4	61.2
16		Order size	50.0	57.7	150.0	500.0
17		Time between orders	0.10	0.12	0.30	1.00
18		Variable cost	€ 2,750.00	€ 2,598.08	€ 3,687.50	€ 10,150.00
19		Unit cost	€ 125,000.00	€ 112,500.00	€ 106,250.00	€ 100,000.00
20		Total cost	€ 127,750.00	€ 115,098.08	€ 109,937.50	€ 110,150.00
21					**minimum**	

Figure 4.8 Best order size with discounts in unit price

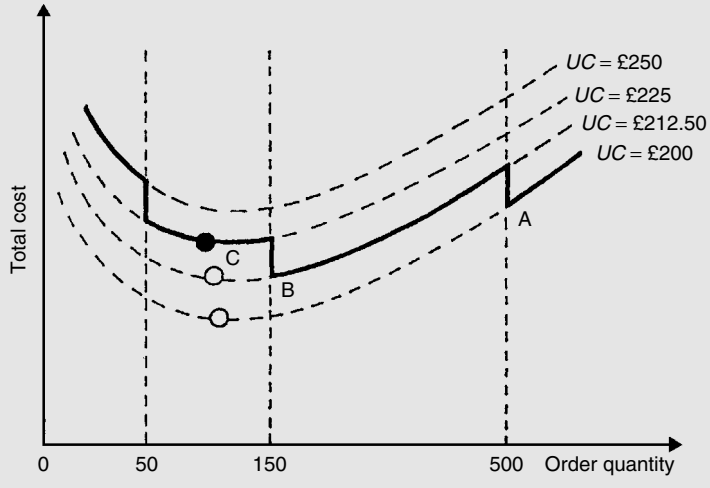


Figure 4.9 Valid cost curve for worked example

Rising delivery cost

We can extend this method for considering unit costs that fall in discrete steps to any problem where there is a discrete change in cost. Suppose, for example, the cost of delivery – and therefore the reorder cost – rises in steps as more vehicles are needed. We can use the same basic approach, as shown in the following example.

Worked example

Kwok Cheng Ho makes a range of high quality garden ornaments. On an average day he uses 4 tonnes of fine grain sand. This sand costs £20 a tonne to buy, and £1.90 a tonne to store for a day. Deliveries are made by modified lorries that carry up to 15 tonnes, and each delivery of a load or part load costs £200. Find the cheapest way to ensure continuous supplies of sand.

Solution

Here the unit cost is constant, but the reorder cost is £200 for each lorry load. This means that orders up to 15 tonnes need one lorry with a reorder cost of £200, orders between 15 and 30 tonnes need two lorries with a reorder costs of £400, orders between 30 and 45 tonnes need three lorries with reorder costs of £600, and so on. Other variables are:

$$D = 4 \text{ tonnes a day}$$

$$UC = £20 \text{ a tonne}$$

$$HC = £1.90 \text{ a tonne a day}$$

We can, again, plot a family of total cost curves. This has the same general form as our price discount curves, but this time the price rises with order quantity and we get the result shown in Figure 4.10 (as before, the continuous line shows the valid cost curve while dotted lines show invalid sections). We want to find the lowest point on the valid curve. Again this lowest point must be at either a

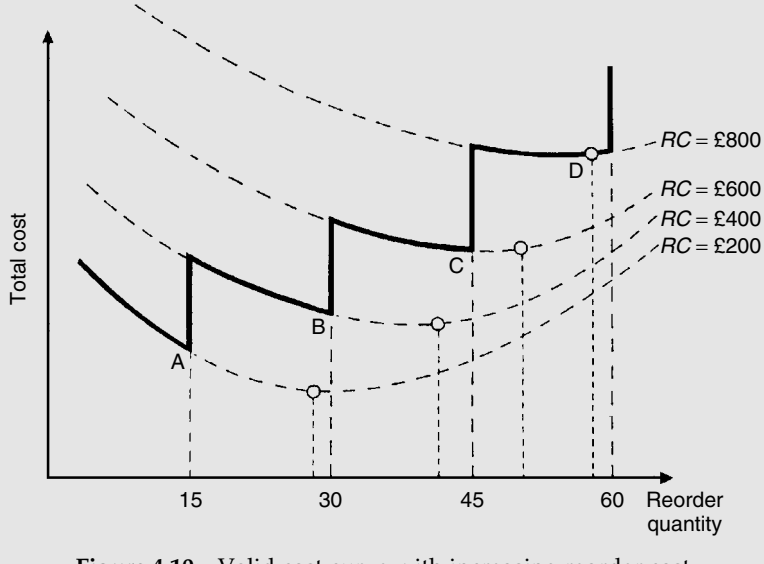


Figure 4.10 Valid cost curve with increasing reorder cost

valid minimum, or at a cost break-point, so we can follow the same procedure as before. The only difference is that as the valid cost curve is reversed and rises to the right, so we look for optimal order quantities to the *left* of a valid minimum and calculate the cost of break-points at the upper limit of valid ranges.

Taking the lowest cost curve:

- RC = £200, valid for Q less than 15 tonnes
- Qo = $\sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 200 \times 4/1.90)} = 29.0$ tonnes
- This is an *invalid* minimum as Qo is not less than 15 tonnes.
- Calculate the cost at the break-point (in this case you can see from Figure 4.10 that the lowest cost is to the right of the range at the higher break-point of 15 tonnes):

$$\begin{aligned} TC &= UC \times D + (RC \times D)/Q + (HC \times Q)/2 \\ &= 20 \times 4 + (200 \times 4)/15 + (1.9 \times 15)/2 \\ &= £147.58 \text{ a day(point A in Figure 4.10).} \end{aligned}$$

Taking the next lowest cost curve:

- RC = £400, valid for Q between 15 and 30 tonnes
- Qo = $\sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 400 \times 4/1.9)} = 41.0$ tonnes
- This is an *invalid* minimum as Qo is not between 15 and 30 tonnes.
- Calculate the cost at the cost break-point (again to the right of the range)

$$\begin{aligned} TC &= UC \times D + (RC \times D)/Q + (HC \times Q)/2 \\ &= 20 \times 4 + (400 \times 4)/30 + (1.9 \times 30)/2 \\ &= £161.83 \text{ a day(point B in Figure 4.10).} \end{aligned}$$

Taking the next lowest cost curve:

- RC = 600 valid for Q between 30 and 45 tonnes
- Qo = $\sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 600 \times 4/1.9)} = 50.3$ tonnes
- This is an *invalid* minimum as Qo is not between 30 and 45 tonnes.
- Calculating the cost at the break-point (again to the right of the range):

$$\begin{aligned} TC &= UC \times D + (RC \times D)/Q + (HC \times Q)/2 \\ &= 20 \times 4 + (600 \times 4)/45 + (1.9 \times 45)/2 \\ &= £176.08 \text{ a day(point C in Figure 4.10).} \end{aligned}$$

Taking the next lowest cost curve:

- RC = £800, valid for Q between 45 and 60 tonnes
- Qo = $\sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 800 \times 4/1.9)} = 58.0$ tonnes
- This is a *valid* minimum as Qo is between 45 and 60 tonnes.
- Calculating the cost at this valid minimum:

$$\begin{aligned} TC &= UC \times D + HC \times Qo \\ &= 20 \times 4 + 1.9 \times 58 \\ &= £190.20 \text{ a day (point D in Figure 4.10).} \end{aligned}$$

As a valid minimum has been reached we stop the procedure, compare costs and find the lowest.

point A	Q = 15 tonnes	cost = £147.58 a day
point B	Q = 30 tonnes	cost = £161.83 a day
point C	Q = 45 tonnes	cost = £176.08 a day
point D	Q = 58 tonnes	cost = £190.20 a day

The best option is to order 15 tonnes at a time, with deliveries needed every 15/4 = 3.75 days.

Summary

Costs often vary with the quantity ordered. A common example has step reductions in unit cost for larger orders. For such problems we can draw a family of curves and find the valid total cost curve. The lowest point on this is either at a valid minimum, or a cost break-point. There is a standard procedure to find this minimum point and identify the best order size. We can adjust this standard procedure to deal with other situations, such as a discrete increase in reorder cost.

Review questions

4.1 Which type of cost can vary with order quantity?

4.2 What is a 'valid minimum' on a total cost curve?

4.3 Where is the minimum point on a valid cost curve?

4.4 If you find a valid minimum on a total cost curve, this shows the best order size. Do you think that this is true?

Finite replenishment rate

Stock from production

You can imagine the economic order quantity working with a wholesaler or retailer. They have large deliveries that instantaneously raise the stock level, and then a series of smaller demands that slowly reduce it. Consider, though, the stock of finished goods at the end of a production line. If the rate of production is greater than demand, goods will accumulate at a finite rate – so there is not instantaneous replenishment, but a finite replenishment rate.

If the rate of production is less than demand, each unit arriving is immediately transferred to a customer and there are no stocks. Goods only accumulate when the production rate is greater than demand. Then the stock level rises at a rate that is the difference between production and demand. If we call the rate of production P, stocks will build up at a rate P − D, as shown in Figure 4.11. This increase will continue as long as production continues. This means that we have to make a decision at some point to stop production of this item – and presumably transfer facilities to making other items. The purpose of this analysis is to find the best time for this transfer, which is equivalent to finding an optimal batch size.

We will suppose that the other assumptions we made for the economic order quantity still hold, so we still consider a single item, with demand that is known, constant and continuous, with costs that are known and constant, and no shortages allowed. So we have replenishment at a rate P and demand at a rate D, with stock growing at a rate P − D. After some time, PT, we decide to stop production. Then, stock is used to meet demand and declines at a rate D. After some further time, DT, all stock has been used and we must start production again. Figure 4.12 shows the resulting variation in stock level, where we assume there is an optimal value for PT (corresponding to an optimal batch size) that we always use.

Optimal batch size

The overall approach of this analysis is the same as the economic order quantity, so we are going to find the total cost for one stock cycle, divide this total cost by the cycle length to get a cost per unit time, and then minimize the cost per unit time.

Consider one cycle of the modified saw-tooth pattern shown in Figure 4.13. If we make a batch of size Q, the maximum stock level with instantaneous replenishment would also be Q. With a finite replenishment rate, though, stock never reaches this

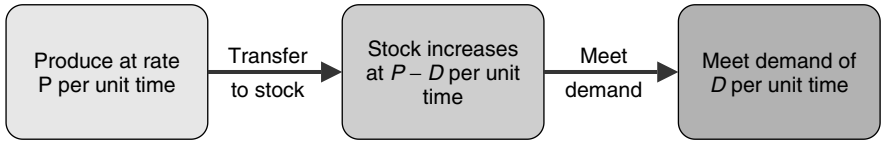

Produce at rate P per unit time
Transfer to stock
Stock increases at P – D per unit time
Meet demand
Meet demand of D per unit time

Figure 4.11 Finite replenishment rate

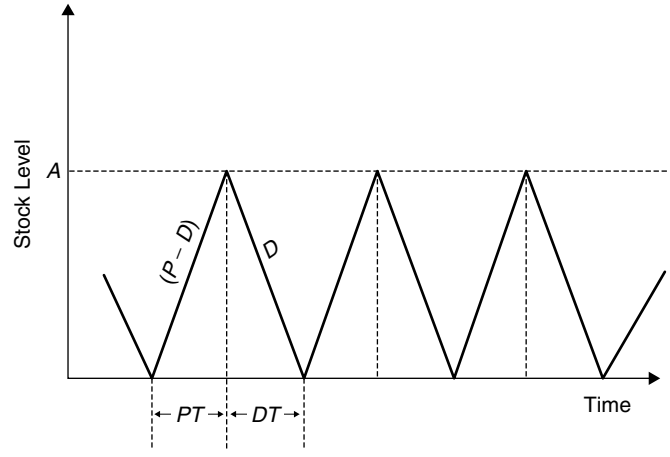


Figure 4.12 Variation in stock level with finite replenishment rate

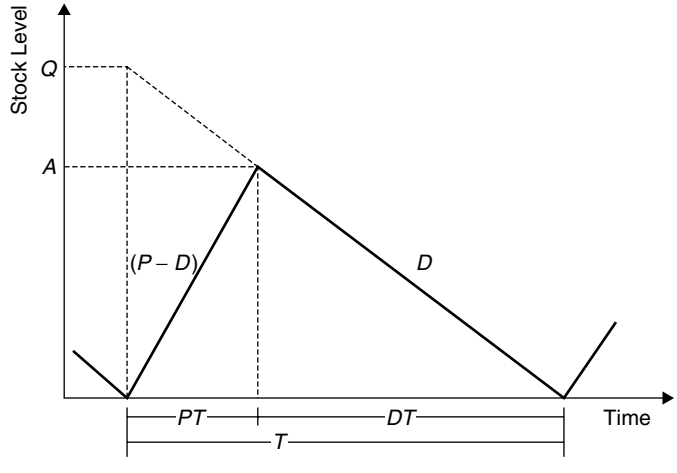


Figure 4.13 One stock cycle with a finite replenishment rate

level, as units are continuously being removed to meet demand. The maximum stock level is lower than Q and occurs at the point where production stops.

Looking at the productive part of the cycle we have:

$$A = (P - D) \times PT$$

We also know that total production during the period is:

$$Q = P \times PT \quad \text{or} \quad PT = Q/P$$

Substituting for PT into the equation for A gives:

$$A = Q \times \frac{(P - D)}{P}$$

Now we can add the three cost components (unit, reorder and holding) for a cycle, remembering that the reorder cost may really be a production set-up cost.

- unit cost component : = number of units made (Q) × unit (UC)
$$= UC \times Q$$
- reorder cost component : = number of production set-ups (1)
$$\times \text{ reorder cost (RC)} = RC$$
- holding cost component : = average stock level (A/2) × time held (T)
$$\times \text{ holding cost (HC)}$$
$$= \frac{HC \times A \times T}{2} = \frac{HC \times Q \times T}{2} \times \frac{(P - D)}{P}$$

Adding these three components gives the total cost for the cycle as:

$$UC \times Q + RC + \frac{HC \times Q \times T}{2} \times \frac{(P - D)}{P}$$

Dividing this by the cycle length, T, gives the total cost per unit time, TC:

$$TC = \frac{UC \times Q}{T} + \frac{RC}{T} + \frac{HC \times Q}{2} \times \frac{(P - D)}{P}$$

Then substitution of $Q = D \times T$, or $T = Q/D$, gives:

$$TC = UC \times D + \frac{RC \times D}{Q} + \frac{HC \times Q}{2} \times \frac{(P - D)}{P}$$

If you compare this with the economic order quantity calculation, the only difference is the factor $(P - D)/P$. We could, therefore, draw a graph of this total cost against batch size and again get an asymmetric 'U'-shaped curve with a distinct minimum. We can find this minimum by differentiating TC with respect to Q and equating the result to zero.

$$\frac{d(TC)}{dQ} = -\frac{RC \times D}{Qo^2} + \frac{HC}{2} \times \left[\frac{P - D}{P}\right] = 0$$

or

$$Qo^2 = \left[\frac{2 \times RC \times D}{HC}\right] \times \left[\frac{P}{P - D}\right]$$

Then

$$\boxed{\text{Optimal batch size} = Qo = \sqrt{\frac{2 \times RC \times D}{HC}} \times \sqrt{\frac{P}{P - D}}}$$

Table 4.1 Comparison of finite and instantaneous replenishment

	Finite replenishment rate	Instantaneous replenishment (for the economic order quantity)
Order quantity	$Qo = \sqrt{\frac{2 \times RC \times D}{HC}} \times \sqrt{\frac{P}{P - D}}$	$Qo = \sqrt{\frac{2 \times RC \times D}{HC}}$
Cycle length	$To = \sqrt{\frac{2 \times RC}{HC \times D}} \times \sqrt{\frac{P}{P - D}}$	$To = \sqrt{\frac{2 \times RC}{HC \times D}}$
Variable cost	$VCo = \sqrt{2 \times RC \times HC \times D} \times \sqrt{\frac{P - D}{P}}$	$VCo = \sqrt{2 \times RC \times HC \times D}$
Total cost	$TCo = UC \times D + VCo$	$TCo = UC \times D + VCo$
Production time	$PTo = Qo/P$	

Now we can substitute this value for Qo into $Qo = D \times To$ and find the best time between orders, and substitute it into $Qo = P \times PT$ to get the best length of production. These results are summarized in Table 4.1. As you can see, they only differ from the results for the economic order quantity by the factor $\sqrt{(P/(P - D))}$. The results again show the compromise between large, infrequent batches (with high holding cost component) and small, frequent batches (with high reorder cost component). With a finite replenishment rate the stock level is somewhat lower than it would be with instantaneous replenishment, so we would expect – all other things being equal – to make larger batches. You can notice from these results that we can no longer assume that the variable cost equals $HC \times Qo$, but have to use the full calculation.

Worked example

Demand for an item is constant at 1,800 units a year. The item can be made at a constant rate of 3,500 units a year. Unit cost is £50, batch set-up cost is £650, and holding cost is 30 per cent of value a year. What is the optimal batch size for the item? If production set-up time is 2 weeks, when should this be started?

Solution

Listing the variables we know:

$$D = 1{,}800 \text{ units a year}$$
$$P = 3{,}500 \text{ units a year}$$
$$UC = £50 \text{ a unit}$$
$$RC = £650 \text{ a batch}$$
$$HC = 0.3 \times 50 = £15 \text{ a unit a year}$$

Substituting these values gives an optimal batch size, Qo, of:

$$Qo = \sqrt{\frac{2 \times RC \times D}{HC}} \times \sqrt{\frac{P}{P-D}} = \sqrt{\frac{2 \times 650 \times 1{,}800}{15}} \times \sqrt{\frac{3{,}500}{3{,}500-1{,}800}} = 566.7$$

Then the optimal production time, PTo, is:

$$PTo = Qo/P = 566.7/3{,}500 = 0.16 \text{ years or } 8.4 \text{ weeks}$$

And the optimal cycle length, To, is:

$$To = \sqrt{\frac{2 \times RC}{HC \times D}} \times \sqrt{\frac{P}{P-D}} = \sqrt{\frac{2 \times 650}{15 \times 1{,}800}} \times \sqrt{\frac{3{,}500}{3{,}500-1{,}800}}$$
$$= 0.31 \text{ years or } 16.4 \text{ weeks}$$

The optimal variable cost, VCo, is:

$$VCo = \sqrt{2 \times RC \times HC \times D} \times \sqrt{\frac{P-D}{P}} = \sqrt{2 \times 650 \times 15 \times 1{,}800}$$
$$\times \sqrt{\frac{3{,}500-1{,}800}{3{,}500}} = £4{,}129 \text{ a year}$$

The optimal total cost per unit time, Tco, is:

$$TCo = UC \times D + VCo = 50 \times 1{,}800 + 4{,}129 = £94{,}129 \text{ a year.}$$

If production set-up takes 2 weeks, we can find the time to start from the reorder level calculation. The cycle time is longer than lead time, so:

$$ROL = LT \times D = 2 \times (1{,}800/52) = 70 \text{ (rounding up to the nearest integer)}$$

The best policy is to start making a batch of 567 units whenever stocks fall to 70 units.

Worked example

Forecasters at Saloman Curtis have estimated the demand for an item that they import from Taiwan to average 20 units a month. They pay €1,000 for each unit. Last year the Purchasing and Inward Transport Department arranged the delivery of 2,000 orders and had running costs of €5,000,000. The Accounts Section quotes the annual holding costs as 20 per cent of unit cost for capital and opportunity loss, 5 per cent for storage space, 3 per cent for deterioration and obsolescence and 2 per cent for insurance. All other costs associated with

stocking the item are combined into a fixed annual cost of €24,000. Calculate the economic order quantity for the item, the time between orders and the corresponding total cost.

By making the item at a rate of 40 units a month Saloman Curtis could avoid the fixed cost of €24,000 a year, reduce the unit cost to €900 and have a batch set-up cost of €1,000. Would it be better for the company to make the item itself rather than buy it in?

Solution

First, we can consider the option of buying the item. Then we know the following variables.

$$D = 20 \times 12 = 240 \text{ units a year}$$

$$UC = €1{,}000 \text{ a unit}$$

$$RC = 5{,}000{,}000/2{,}000 = €2{,}500 \text{ an order}$$

$$HC = (0.2 + 0.05 + 0.03 + 0.02) \times 1{,}000 = 0.3 \times 100 = €300 \text{ a unit a year}$$

It is easiest to do these calculations on a spreadsheet, and Figure 4.14 shows a summary of results. As you can see, the total cost of buying the item is €258,974 a year. We have to add to this the fixed charges of €24,000 to get a total of €282,974 a year.

	A	B	C	D	E	F	G	H	I
1	**Inventory Control – Economic Order Quantity**								
2									
3	**Notes:**	Calculations for the economic order quantity,							
4		including cost sensitivity.							
5		All units are consistent, so a lead time of one							
6		week appears as 1/52 years.							
7									
8	**Inputs**	**Demand**		240		**Variable cost range**			
9		**Unit cost**	€	1,000.00		**Percentage**	10%		
10		**Holding cost amount**	€	-		**Proportion**	1.1		
11		**Holding cost percentage**		30%					
12		**Reorder cost**	€	2,500.00					
13		**Lead time**		0					
14									
15	**Results**	**Holding cost**	€	300.00		**Order quantity range**			
16	**Orders**	**Order size**		63.25		**Proportion, f**	0.64	to	1.56
17		**Time between orders**		0.26		**Order size**	40.59	to	98.55
18		**Number of orders a period**		3.79					
19		**Reorder level**		0.00					
20	**Investment**	**Average stock**		31.62					
21		**Average investment**	€	31,622.78					
22		**Minimum stock**		0					
23		**Minimum investment**	€	-					
24		**Maximum stock**		63.25					
25		**Maximum investment**	€	63,245.55					
26	**Costs**	**Fixed unit costs**	€	240,000.00					
27		**Variable cost**	€	18,973.67					
28		**Total cost a period**	€	258,973.67					

Figure 4.14 The cost of buying the item from Taiwan

Alternatively, Saloman Curtis could make the item itself with:

$$P = 40 \times 12 = 480 \text{ units a year}$$

$$UC = €900 \text{ a unit}$$

$$RC = €1{,}000 \text{ a batch}$$

$$HC = 0.3 \times 900 = €270 \text{ a unit a year}$$

Again, we can use a spreadsheet for the calculations, as illustrated in Figure 4.15. This shows that the cost of making the item is €224,050 a year, a saving of almost €60,000 a year. Based on these figures it seems attractive for Saloman Curtis to make the item themselves, but there are a lot of other factors to consider.

	A	B	C	D	E	F
1	**Inventory Control – best order quantity with finite production rate**					
2						
3	**Notes:**	Calculations for the economic order quantity,				
4		with production at rate P.				
5		Now Qo = sqrt((2*RC*D)/HC)*sqrt((P)/(P-D))				
6		and other calculations are adjusted accordingly.				
7						
8	**Inputs**	Demand		240		
9		Unit cost		€ 900.00		
10		Holding cost amount		€ -		
11		Holding cost percentage		30.0%		
12		Reorder cost		€ 1,000.00		
13		Lead time				
14		Production rate		480		
15						
16	**Results**	Holding cost		€ 270.00		
17	**Orders**	Order size		59.63		
18		Time between orders		0.248		
19		Number of orders a period		4.02		
20		Reorder level				
21	**Investment**	Average stock		29.81		
22		Average investment		€ 26,832.82		
23		Minimum stock		0		
24		Minimum investment		€ -		
25		Maximum stock		59.63		
26		Maximum investment		€ 53,665.63		
27	**Costs**	Fixed unit costs		€ 216,000.00		
28		Variable cost		€ 8,049.84		
29		Total cost a period		€ 224,049.84		

Figure 4.15 Finding the best policy for making the item

Summary

Replenishment of stock often occurs at a finite rate rather than instantaneously. This gives a cycle with stock levels rising at a rate $P - D$ during replenishment and then falling at a rate D when replenishment stops. We can find an optimal batch size. This, and related results, vary from the economic order calculation by the factor $\sqrt{(P/(P - D))}$.

Review questions

4.5 What is the maximum stock level when demand is greater than the replenishment rate?

4.6 Is DT (the part of a stock cycle when there is no production) longer or shorter than PT (the part when there is production)?

4.7 Does a finite replenishment rate lead to larger or smaller batches than the economic order quantity?

4.8 What happens to this analysis when the demand is greater than the replenishment rate?

Planned shortages with back-orders

Back-orders and lost sales

The models described so far have assumed that no shortages are allowed and that all demand must be met. This is a reasonable view when shortages are very expensive. There are, however, circumstances where planned shortages are beneficial. An obvious example comes when the cost of keeping an item in stock is higher than the profit from selling it. When you buy a new car, for example, the showroom does not keep a stock of every variation on its models, but you choose the features you want, the showroom orders this from the manufacturer, and then you wait for the car to be delivered. If you buy a computer from Dell, you visit their website to design a systems, and then Dell make this and deliver it with a short lead time.

When customer demand for an item cannot be met from stock, there are shortages. Then customer have a choice (shown in Figure 4.16). First, they can wait for the item to come into stock, in which case their demand is met by a *back-order*. Or they can withdraw their order and go to another supplier, in which case there are *lost sales*. Customers who experience a shortage are likely to divert at least some future business to more reliable suppliers. Surveys of customer attitudes suggest that it takes one poor experience for a customer to change supplier, but something like 14 good experiences to restore their confidence. We will start by looking at back-orders, where customers are prepared to wait, and then move on to lost sales.

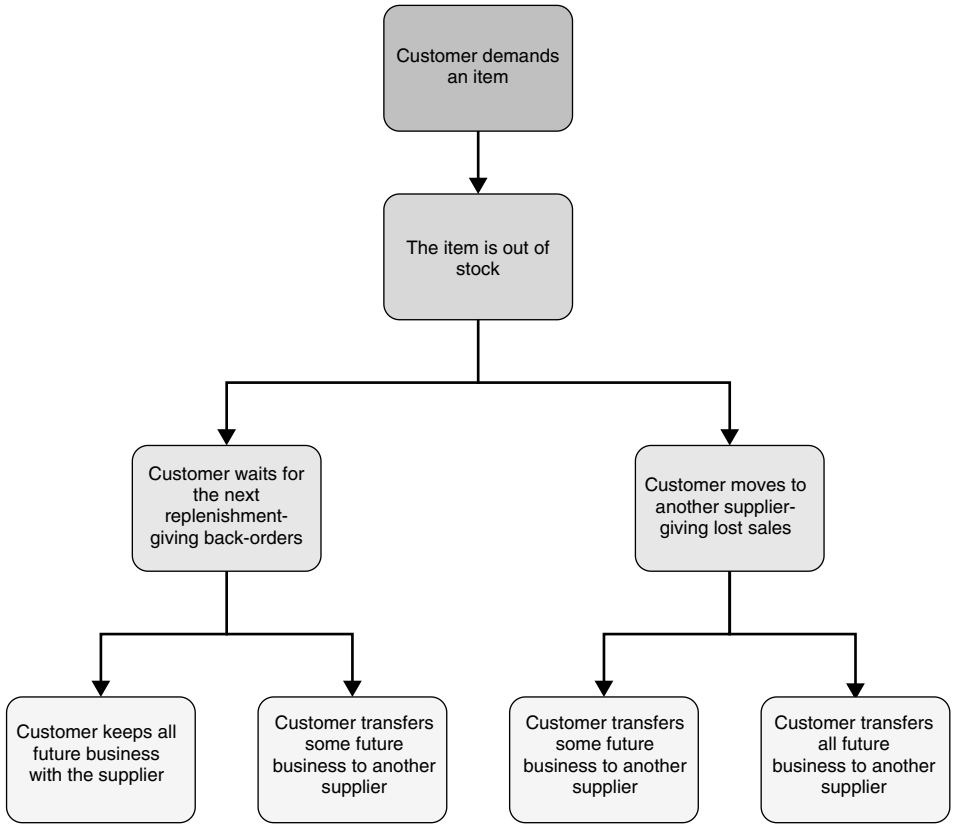

Customer demands an item
The item is out of stock
Customer waits for the next replenishment-giving back-orders
Customer moves to another supplier-giving lost sales
Customer keeps all future business with the supplier
Customer transfers some future business to another supplier
Customer transfers some future business to another supplier
Customer transfers all future business to another supplier

Figure 4.16 Alternatives for customers when their demand cannot be met from stock

Back-orders

A back-order occurs when a customer demands an item that is out of stock, and then waits to receive the item from the next delivery from suppliers. You see this in many retailers, such as furniture showrooms. Each showroom stocks a selection of furniture, but there is not enough to cover all demand, and customers are asked to wait for deliveries from suppliers or regional distribution centres. This suggests that back-ordering is more likely when the unit cost is high, there is a wide range of items, it would be too expensive to hold stocks of the complete range, lead times from suppliers are reasonably short, there is limited competition, and customers are prepared to wait.

In extreme cases an organization keeps no stock at all and meets all demand by back orders. This gives 'make to order' rather than 'make for stock' operations. We will look at the more common case where some stock is kept, but not enough to cover all demand. The key question is, how much of the demand should be met from stock and how much from back-orders?

There is always some cost associated with back-orders, including administration, loss of goodwill, some loss of future orders, emergency orders, expediting, and so on. This cost is likely to rise with increasing delay. We can, then, define a shortage cost, SC, which is time-dependent and is a cost per unit per unit time delayed.

Figure 4.17 shows the stock level during one cycle when shortages are back-ordered. Here back-orders are shown as negative stocks, and we are going to use the standard approach of finding the cost for a single cycle and using this to calculate the optimal order size.

The total cost for a single cycle comes from adding the four cost components:

- unit cost component: unit cost time number of units bought $= UC \times Q$
- reorder cost component: reorder cost times number of orders $= RC$
- holding cost component: an average stock of $(Q - S)/2$ held for a time T_1
$$= \frac{HC \times (Q - S) \times T_1}{2}$$
- shortage cost component: an average shortage of $S/2$ held for a time T_2
$$= \frac{SC \times S \times T_2}{2}$$

Adding these together gives the total cost per cycle:

$$UC \times Q + RC + \frac{HC \times (Q - S) \times T_1}{2} + \frac{SC \times S \times T_2}{2}$$

During the first part of the cycle all demand is met from stock, so the amount sent to customers is $Q - S$, which equals the demand of $D \times T_1$. During the second

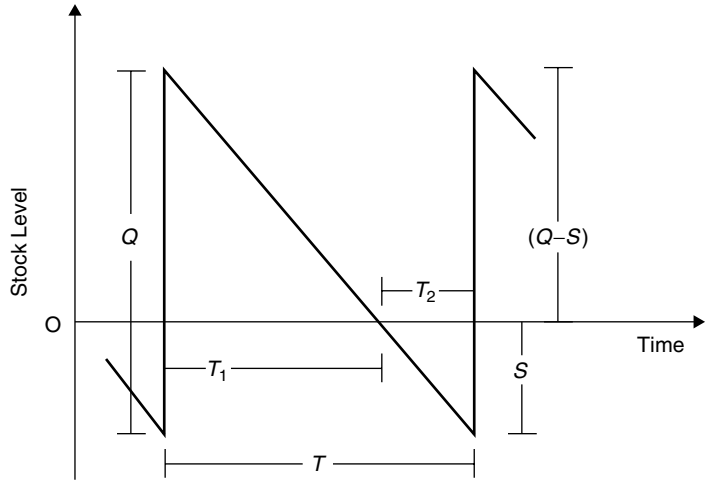


Figure 4.17 A single stock cycle with back-orders

part of the cycle all demand is back-ordered, so the shortage, S equals the unmet demand of $D \times T_2$. Substituting:

$$T_1 = (Q - S)/D \quad \text{and} \quad T_2 = S/D$$

gives the total cost per cycle as:

$$UC \times Q + RC + \frac{HC \times (Q - S)^2}{2 \times D} + \frac{SC \times S^2}{2 \times D}$$

Then dividing by T and substituting $Q = D \times T$ gives the total cost per unit time:

$$TC = UC \times D + \frac{RC \times D}{Q} + \frac{HC \times (Q - S)^2}{2 \times Q} + \frac{SC \times S^2}{2 \times Q}$$

The equation has two variables, Q and S, so we can differentiate with respect to both of these and set the results to zero:

$$\frac{\delta(TC)}{\delta Q} = 0 = -\frac{RC \times D}{Q^2} + \frac{HC}{2} - \frac{HC \times S^2}{2 \times Q^2} - \frac{SC \times S^2}{2 \times Q^2}$$

and

$$\frac{\delta(TC)}{\delta S} = 0 = -HC + \frac{HC \times S}{Q} - \frac{SC \times S}{Q}$$

After some manipulation these simultaneous equations can be solved to give optimal values:

Optimal order size

$$Qo = \sqrt{\frac{2 \times RC \times D \times (HC + SC)}{HC \times SC}}$$

Optimal amount to be back-ordered

$$To = \sqrt{\frac{2 \times RC \times HC \times D}{SC \times (HC + SC)}}$$

In addition we know:

$T_1 = (Qo - So)/D$ = time during which demand is met

$T_2 = So/D$ = time during which demand is back-ordered

$T = T_1 + T_2$ = cycle time

Worked example

Demand for an item is constant at 100 units a month. Unit cost is £50, reorder cost is £50, holding cost is 25 per cent of value a year, shortage cost for back-orders is 40 per cent of value a year. Find an optimal inventory policy for the item.

Solution

Listing the variables we know in consistent units:

$$D = 100 \times 12 = 1{,}200 \text{ units a year}$$

$$UC = £50 \text{ a unit}$$

$$RC = £50 \text{ an order}$$

$$HC = 0.25 \times 50 = £12.5 \text{ a unit a year}$$

$$SC = 0.4 \times 50 = £20 \text{ a unit a year}$$

Then substitution gives:

$$Qo = \sqrt{\frac{2 \times RC \times D \times (HC + SC)}{HC \times SC}} = \sqrt{\frac{2 \times 50 \times 1{,}200 \times (12.5 + 20)}{12.5 \times 20}} = 125 \text{ units}$$

$$So = \sqrt{\frac{2 \times RC \times HC \times D}{SC \times (HC + SC)}} = \sqrt{\frac{2 \times 50 \times 12.5 \times 1{,}200}{20 \times (12.5 + 20)}} = 48 \text{ units}$$

$$T_1 = (Qo - So)/D = (125-48)/1{,}200 = 0.064 \text{ years} = 3.3 \text{ weeks}$$

$$T_2 = So/D = 48/1{,}200 = 0.04 \text{ years} = 2.1 \text{ weeks}$$

$$T = T_1 + T_2 = 5.4 \text{ weeks}$$

In this example the shortage cost is relatively low, so the item is out of stock almost 40 per cent of the time.

Summary

Sometimes it is useful to have planned shortages, particularly when customer demand is not lost but can be met from back-orders. These back-orders inevitably have some cost, which we can express as a time-dependent shortage cost. We can use this to find optimal values for the order size and amount to be back-ordered.

Review questions

4.9 When are inventory items back-ordered?

4.10 What do you think is the main problem with using the back-order analysis described?

4.11 Are inventory systems with back-orders always more expensive than those without back-orders?

Lost sales

When there are shortages, customers may not be willing to wait for back-orders to arrive, but will simply move to another supplier. If a newsagent has run out of a magazine that you want, you do not wait for the next delivery but simply go to another seller down the road. This gives lost sales, and the pattern of stock shown in Figure 4.18.

An initial stock of Q runs out after a time Q/D, and all subsequent demand is lost until the next replenishment arrives. We can no longer say that $Q = D \times T$ as there is unsatisfied demand and the amount supplied in a cycle is less than demand. In particular, there is an unsatisfied demand of $D \times T - Q$.

You can also see that when shortages are allowed, the aim of minimizing cost is no longer the same as maximizing revenue. If shortages are allowed we might, in some circumstances, minimize costs by holding no stock at all – but this would certainly not maximize revenue. In this analysis we will maximize the net revenue, which is defined as the gross revenue minus costs. For this we have to define SP as the selling price per unit. We also need to look at the cost of lost sales, which has two parts. First, there is a loss of profit, which is a notional cost that we can define as SP − UC per unit of sales lost. Second, there is a direct cost, which includes loss of goodwill, remedial action, emergency procedures, and so on. We will define this as DC per unit of sales lost. Now we can use these in our usual approach.

The four cost components for a single stock cycle are:

- unit cost component: $UC \times Q$
- reorder cost component: RC
- holding cost component: an average stock of Q/2 held for time Q/D
  $$= \frac{HC \times Q}{2} \times \frac{Q}{D}$$
- lost sales cost component: taking only the actual cost of DC for each of $D \times T - Q$ lost sales $= DC \times (D \times T - Q)$

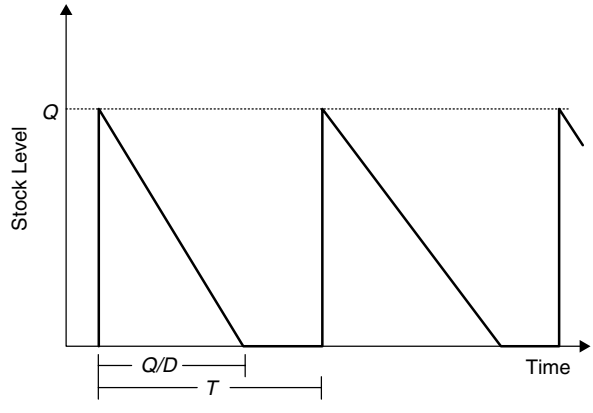


Figure 4.18 Stock level with lost sales

Then the net revenue per cycle is the gross revenue (SP × Q) minus the sum of these costs:

$$SP \times Q - UC \times Q - RC - \frac{HC \times Q^2}{2 \times D} - DC \times (D \times T - Q)$$

Dividing this by T gives the net revenue per unit time:

$$R = \frac{1}{T} \times \left[Q \times (DC + SP - UC) - RC - \frac{HC \times Q^2}{2 \times D} - DC \times D \times T\right]$$

Now we can define the cost of each unit of lost sales *including* loss of profit:

$$\begin{aligned} LC &= \text{cost of each unit of lost sales including loss of profits} \\ &= DC + SP - UC \end{aligned}$$

and we can set:

$$\begin{aligned} Z &= \text{proportion of demand satisfied} \\ &= Q/(D \times T) \end{aligned}$$

Substituting these into the equation for R gives:

$$R = Z \times \left[D \times LC - \frac{RC \times D}{Q} - \frac{HC \times Q}{2} - DC \times D\right]$$

Differentiating this with respect to Q and setting the result to zero gives a maximum value for net revenue:

$$\frac{dR}{dQ} = 0 = -\frac{Z \times RC \times D}{Q^2} - \frac{Z \times HC}{2}$$

or

$$Qo = \sqrt{\frac{2 \times RC \times D}{HC}}$$

which is the standard economic order quantity. When we substitute this into the revenue equation we get the optimal value for R, which we will call Ro:

$$Ro = Z \times [D \times LC - \sqrt{2 \times RC \times HC \times D}]$$

Now Z is the proportion of demand met, so we can set this to any value in the range zero to one. But we want to choose the value that maximizes the revenue, Ro. We find this using the following argument.

- If D × LC is greater than $\sqrt{(2 \times RC \times HC \times D)}$ the term in brackets is positive, the revenue is positive and want to make Z as large as possible. So Z = 1 and there are no shortages.
- If D × LC is less than $\sqrt{(2 \times RC \times HC \times D)}$ the term in brackets is negative, we make a loss and we want to make Z as small as possible. So Z = 0 and we do not stock the item at all.

- If $D \times LC$ equals $\sqrt{(2 \times RC \times HC \times D)}$ the term in brackets equals zero the revenue is zero whatever value we give to Z.

Worked example

Find the best ordering policy for three items with the following costs.

Item	D	RC	HC	DC	SP	UC
1	50	150	80	20	110	90
2	100	400	200	10	200	170
3	50	500	400	30	350	320

Solution

For each item we have to calculate $D \times LC$ and compare this with $\sqrt{(2 \times RC \times HC \times D)}$. We also have to remember that, by definition, $LC = DC + SP - UC$.

Item 1

$$LC = DC + SP - UC = 20 + 110 - 90 = 40$$

$$D \times LC = 50 \times 40 = 2{,}000$$

$$\sqrt{(2 \times RC \times HC \times D)} = \sqrt{(2 \times 150 \times 80 \times 50)} = 1{,}095$$

As $D \times LC$ is greater than $\sqrt{(2 \times RC \times HC \times D)}$ the net revenue is positive and we set $Z = 1$. All demand is met and no sales are lost. Then:

$$Ro = Z \times [D \times LC - \sqrt{(2 \times RC \times HC \times D)}] = 1 \times [2{,}000 - 1{,}095] = 905$$

$$Qo = \sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 150 \times 50/80)} = 13.7$$

Item 2

$$LC = DC + SP - UC = 10 + 200 - 170 = 40$$

$$D \times LC = 100 \times 40 = 4{,}000$$

$$\sqrt{(2 \times RC \times HC \times D)} = \sqrt{(2 \times 400 \times 200 \times 100)} = 4{,}000$$

As $D \times LC = \sqrt{(2 \times RC \times HC \times D)}$ the net revenue is zero and we can set Z to any convenient value. Then:

$$Ro = 0$$

$$Qo = \sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 400 \times 100/20)} = 200$$

Item 3

$$LC = DC + SP - UC = 30 + 350 - 320 = 60$$

$$D \times LC = 50 \times 60 = 3{,}000$$

$$\sqrt{(2 \times RC \times HC \times D)} = \sqrt{(2 \times 500 \times 400 \times 50)} = 4{,}472$$

As $D \times DC$ is less than $\sqrt{(2 \times RC \times HC \times D)}$ the net revenue is negative and we set $Z = 0$. The item is not stocked at all.

Summary

Shortages often lead to lost sales rather than back-orders. We can build a model for shortages that maximizes net revenues, rather than minimizing total costs. This gives a simple rule that shows whether or not to stock an item.

Review questions

4.12 When are sales lost?

4.13 What does it mean when $D \times LC = \sqrt{(2 \times RC \times HC \times D)}$?

4.14 Why does the analysis for lost sales maximize revenue rather than minimize costs?

4.15 How would the economic order quantity change if we maximize revenue rather than minimize cost?

Constraints on stock

So far we have assumed that each inventory item is completely independent. Then we can calculate an optimal order policy for each item in isolation and there is no need to consider interactions with other items. In practice there are situations where, although demand for each item is independent, we may have to look at these interactions. An obvious example happens when several items are ordered from the same supplier; and we can reduce delivery costs by combining orders for different items in a single delivery. Another example happens when there are constraints on some operations, such as limited warehouse space or a maximum acceptable investment in stock.

Problems of this type can become rather complicated, so we will illustrate some broader principles by looking at problems with constraints on stock levels. These constraints could include limited storage space, a maximum acceptable investment in stock, a maximum number of deliveries that can be accepted, maximum size of delivery that can be handled, a maximum number of orders that can be placed, and so on. You can tackle most problems with constraints using the same methods, so we will describe two of the most useful analyses. The first shows an intuitive

approach to problems where there is a limited amount of storage space; the second is a more formal analysis for constrained investment.

Constraints on storage space

If an organization uses the economic order quantity for all items in an inventory, the resulting total stock might exceed the available capacity. Then we need some way of reducing the stock until it is within acceptable limits. One approach puts an additional cost on space used. Then the holding cost is in two parts: the original holding cost, HC, which we have used before and an additional cost, AC, related to the storage area (or volume) used by each unit of the item. Then the total holding cost per unit per unit time becomes $HC + AC \times S_i$, where S_i is the amount of space occupied by one unit of item i. When we use this revised cost in the economic order quantity calculation we find that:

$$Q_i = \sqrt{\frac{2 \times RC_i \times D_i}{HC_i + AC \times S_i}}$$

Here we have used subscripts for all variables to show that they can be different for each item, i.

If there are no constraints on space, we can give AC a value of zero, and the result is the standard economic order quantity. If, however, space is constrained we can give AC a positive value and order quantities of all items are reduced from Qo_i to Q_i. As the average stock level is $Q_i/2$, this automatically reduces the amount of stock. The reduction necessary – and the consequent value of AC – depend on the severity of space constraints. A reasonable way of finding solutions is iteratively to adjust AC until the space required exactly matches available capacity.

Worked example

The Brendell Work Centre is concerned with the storage space occupied by three items that have the following characteristics.

Item	Unit cost ($)	Demand	Space per unit (m^3)
1	100	500	1
2	200	400	2
3	300	200	3

The reorder cost for the items is constant at $1000, and holding cost is 20 per cent of value a year. If the company wants to allocate an average space of 300 cubic metres to the items, what would be the best ordering policy? How much does the constraint on space increase variable inventory cost?

Solution

We can start by calculating the reorder quantities to see if the space constraint is limiting:

Item 1

$$Qo_1 = \sqrt{(2 \times RC \times D_1/HC_1)} = \sqrt{(2 \times 1{,}000 \times 500/20)} = 223.6 \text{ units}$$

$$\text{Average space} = S_1 \times Qo_1/2 = 1 \times 223.6/2 = 111.8 \text{ cubic metres}$$

Item 2

$$Qo_2 = \sqrt{(2 \times RC \times D_2/HC_2)} = \sqrt{(2 \times 1{,}000 \times 400/40)} = 141.4 \text{ units}$$

$$\text{Average space} = S_2 \times Qo_2/2 = 2 \times 141.2/2 = 141.4 \text{ cubic metres}$$

Item 3

$$Qo_3 = \sqrt{(2 \times RC \times D_3/HC_3)} = \sqrt{(2 \times 1{,}000 \times 200/60)} = 81.6 \text{ units}$$

$$\text{Average space} = S_3 \times Qo_3/2 = 3 \times 81.6/2 = 122.4 \text{ cubic metres}$$

Using economic order quantities, the average space occupied is 375.6 cubic metres. This is above the limit of 300 cubic metres set by the company, so stocks must be reduced. For this we set the additional cost for space, AC, to an arbitrary value, say 1, and calculate modified order quantities.

Item 1

$$Q_1 = \sqrt{(2 \times RC \times D_1/(HC_1 + AC \times S_1))}$$

$$= \sqrt{(2 \times 1{,}000 \times 500/(20{+}1 \times 1))} = 218.2 \text{ units}$$

$$\text{Average space} = S_1 \times Qo_1/2 = 1 \times 218.2/2 = 109.1 \text{ cubic metres}$$

Item 2

$$Q_2 = \sqrt{(2 \times RC \times D_2/(HC_2 \times AC \times S_2))}$$

$$= \sqrt{(2 \times 1{,}000 \times 400/(40{+}1 \times 2))} = 138.0 \text{ units}$$

$$\text{Average space} = S_2 \times Qo_2/2 = 2 \times 138.0/2 = 138.0 \text{ cubic metres}$$

Item 3

$$Q_3 = \sqrt{(2 \times RC \times D_3/(HC_3 + AC \times S_3))}$$

$$= \sqrt{(2 \times 1{,}000 \times 200/(60{+}1 \times 3))} = 79.7 \text{ units}$$

$$\text{Average space} = S_3 \times Qo_3/2 = 3 \times 79.7/2 = 119.6 \text{ cubic metres}$$

This gives an average stock of 366.7 cubic metres. Although this is less than before, it is still above the limits set by the company. Now we repeat the calculations, iteratively increasing AC until we get a result that satisfies the constraint. As you can see from the following table, this happens when AC is around 11 (actually 11.37).

AC	Q_1	Q_2	Q_3	Space needed
0	223.6	141.4	81.6	375.6
1	218.2	138.0	79.7	366.7
5	200.6	126.5	73.0	336.0
10	182.6	115.5	66.7	306.8
11	179.6	113.6	65.6	301.8
12	176.8	111.8	64.5	297.0
13	174.0	110.1	63.6	292.5

To make sure that the stock is kept within the space limit, we might actually use a value of AC = 12. Then we can compare the variable costs with those from the economic order quantity.

Item 1

Using economic order quantity:

$$VCo_1 = HC_1 \times Qo_1 = 20 \times 223.6 = \$4{,}472$$

Using the revised order quantity:

$$VC_1 = RC \times D_1/Q_1 + HC_1 \times Q_1/2 = 1{,}000 \times 500/176.8 + 20 \times 176.8/2 = \$4{,}596$$

Item 2

Using economic order quantity:

$$VCo_2 = HC_2 \times Qo_2 = 40 \times 141.4 = \$5{,}656$$

Using the revised order quantity:

$$VC_2 = RC \times D_2/Q_2 + HC_2 \times Q_2/2 = 1{,}000 \times 400/111.8 + 40 \times 111.8/2 = \$5{,}814$$

Item 3

Using economic order quantity:

$$VCo_3 = HC_3 \times Qo_3 = 60 \times 81.6 = \$4{,}896$$

Using the revised order quantity:

$$VC_3 = RC \times D_3/Q_3 + HC_3 \times Q_3/2 = 1{,}000 \times 200/64.5 + 60 \times 64.5/2 = \$5{,}036$$

To meet the space constraint, the total variable cost has risen from (4,472 + 5,656 + 4,896 =)$15,024 to (4,596 + 5,814 + 5,036 =)$15,446, which is a rise of less than 3 per cent.

Constraints on average investment in stock

When there is a constraint on space, we calculated revised order quantities with an additional cost for space occupied. We can use the same approach for other constraints, such as a maximum average investment in stocks. This time, though, rather than use an iterative procedure to find the optimal, we will use a more formal derivation.

Suppose an organization stocks N items and has an upper limit, UL, on the total average investment. The calculated economic order quantity for each item i is Qo_i, but again we need some means of calculating the best, lower amount Q_i which allows for the constraint. The average stock of item i is $Q_i/2$ and the average investment in this item is $UC \times Q_i/2$. The problem then becomes one of constrained optimization – minimize the total variable cost subject to the constraint of an upper limit on average investment. That is:

$$\text{Minimize:} \quad VC = \sum_{i=1}^{N} \frac{RC_i \times D_i}{Q_i} + \frac{HC_i \times Q_i}{2}$$

$$\text{Subject to:} \quad \sum_{i=1}^{N} \frac{UC_i \times Q_i}{2} \leq UL$$

The usual way of solving problems of this type is to eliminate the constraint by introducing a Lagrange multiplier and incorporating this into a revised objective function. In particular, we can adjust the constraint to make it less than or equal to zero, and the result is multiplied by the Lagrange multiplier, LM. This is incorporated in the revised objective function:

$$\text{Minimize:} \sum_{i=1}^{N} \left[\frac{RC_i \times D_i}{Q_i} + \frac{HC_i \times Q_i}{2} \right] + LM \times \left[\sum_{i=1}^{N} \frac{UC_i \times Q_i}{2} - UL \right]$$

To solve this, we differentiate with respect to both Q_i and LM, and set both of these derivatives to zero to give two simultaneous equations. We can then eliminate the Lagrange multiplier to give the overall result:

$$\boxed{Q_i = Qo_i \times \frac{2 \times UL \times HC}{UC \times \sum_{i=1}^{N} VCo_i}}$$

When the holding cost is a fixed proportion of the unit cost, the best order quantity for each item is a fixed proportion of the economic order quantity.

Worked example

A company is concerned with the stock of three items with the following demand and unit cost.

Item	Demand	Unit cost
1	100	100
2	300	60
3	200	40

The items have a reorder cost of £100 and a holding cost of 20 per cent of value a year. What is the best ordering policy if the company wants to limit average investment in these item to £4,000?

Solution

To start this analysis we need to calculate the $\sum$VCo, the total variable cost if economic order quantities are used.

Item 1

$$Qo_1 = \sqrt{(2 \times RC \times D_1/HC_1)} = \sqrt{(2 \times 100 \times 100/20)} = 31.62 \text{ units}$$

$$VCo_1 = HC_1 \times Qo_1 = 20 \times 31.62 = £632.40 \text{ a year}$$

$$\text{Average stock} = Qo_1/2 = 31.62/2 = 15.81 \text{ units}$$

$$\text{Average investment} = UC_1 \times \text{average stock} = 100 \times 15.81 = £1{,}581.00$$

Item 2

$$Qo_2 = \sqrt{(2 \times RC \times D_2/HC_2)} = \sqrt{(2 \times 100 \times 300/12)} = 70.71 \text{ units}$$

$$VCo_2 = (HC_2 \times Qo_2) = 12 \times 70.71 = £848.52 \text{ a year}$$

$$\text{Average stock} = Qo_2/2 = 70.71/2 = 35.36 \text{ units}$$

$$\text{Average investment} = UC_2 \times \text{average stock} = 60 \times 35.36 = £2{,}121.60$$

Item 3

$$Qo_3 = \sqrt{(2 \times RC \times D_3/HC_3)} = \sqrt{(2 \times 100 \times 200/8)} = 70.71 \text{ units}$$

$$VCo_3 = (HC_3 \times Qo_3) = 8 \times 70.71 = £565.68 \text{ a year}$$

$$\text{Average stock} = Qo_3/2 = 70.71/2 = 35.36 \text{ units}$$

$$\text{Average investment} = UC_3 \times \text{average stock} = 40 \times 35.36 = £1{,}414.40$$

Using economic order quantities, average investment in the three items is (1,581.00 + 2,121.60 + 1,414.40 =)£5,117. This is above the limit of £4,000 set by the company, so the order quantities must be reduced.

The sum of variable costs is:

$$\sum \text{VCo} = 632.40 + 848.52 + 565.68 = £2{,}046.60$$

and substituting these values gives the following revised order quantities.

Item 1

$$Q_1 = Qo_1 \times \frac{2 \times UL \times HC}{UC \times \sum_{i=1}^{N} VCo_i} = 31.62 \times \frac{2 \times 4{,}000 \times 20}{100 \times 2{,}046.60} = 24.72$$

Average investment = $UC \times Q_1/2 = 100 \times 12.36 = £1{,}236$.

Item 2

$$Q_2 = Qo_2 \times \frac{2 \times UL \times HC}{UC \times \sum_{i=1}^{N} VCo_i} = 70.71 \times \frac{2 \times 4{,}000 \times 12}{60 \times 2{,}046.60} = 55.28$$

Average investment = $UC \times Q_2/2 = 60 \times 27.64 = £1{,}658.40$.

Item 3

$$Q_3 = Qo_3 \times \frac{2 \times UL \times HC}{UC \times \sum_{i=1}^{N} VCo_i} = 70.71 \times \frac{2 \times 4{,}000 \times 8}{40 \times 2{,}046.60} = 55.28$$

Average investment = $UC \times Q_3/2 = 40 \times 27.64 = £1{,}105.60$.

The total average investment is now (1,236.00 + 1,658.40 + 1,105.60 =)£4,000. As the holding cost was defined as 20 per cent of unit cost, we could have saved some effort by simply multiplying the economic order quantities by $2 \times 4{,}000 \times 0.2/2{,}046.68 = 0.78$.

Summary

There are many circumstances in which stock items cannot be treated in isolation, including constraints on stock levels or investment. When there is a constraint on space, we can find revised order quantities by adding an additional cost for the space used. This reduces order quantities, and hence average stock levels. When

there is a limit on the average investment in stock, we can again find reduced order quantities. When the holding cost is a fraction of unit cost, these become a fixed proportion of the economic order quantities.

Review questions

4.16 With independent demand inventory models, each item is considered as independent, so we do not consider interactions. Do you think this is true?

4.17 How is the average stock level reduced when storage space is limited?

4.18 If the total investment in stock is limited, will the best order quantity for each item be greater or less than the economic order quantity?

Discrete, variable demand

So far we have assumed that demand is continuous – in effect, each individual demand is so small that it contributes to an overall demand that is continuous. Often, however, an organization has to meet demand for an integer number of units every period. A car showroom, electrical retailer, furniture warehouse or computer store, for example, can only sell a discrete number of units. In this section we will look at a model for discrete, variable demand. This might be our computer store, which is going to sell 50 units this week, 45 units next week, 60 units the following week, and so on. If we know exactly what these demands are going to be, we can build a deterministic model and find an optimal ordering policy.

The way to approach this is to assume there is some optimal number of period's demand that we should combine into a single order. If we order less than this, we are making frequent orders and the reorder cost component is high, giving a high overall cost. If we order more than this, we have higher stocks levels and the holding cost component is high, giving a high overall cost. As we can only order discrete numbers of units, the variable cost is also discrete and has the form shown in Figure 4.19. Our object is to find the optimal value of N, the number of period's demand to combine into a single order.

Consider one order for an item, where we buy enough to satisfy demand for the next N periods. The demand is discrete and varies, so we can define the demand in period i of a cycle as D_i, and an order to cover all demand in the next N periods will be:

$$A = \sum_{i=1}^{N} D_i$$

We will assume that this arrives in stock at one time, so the highest actual stock level is A. This is used to meet demand and we will assume that it falls back to zero, as shown in Figure 4.20. We can approximate the average stock level by A/2 and the cost of holding this is $(A/2) \times N \times HC$.

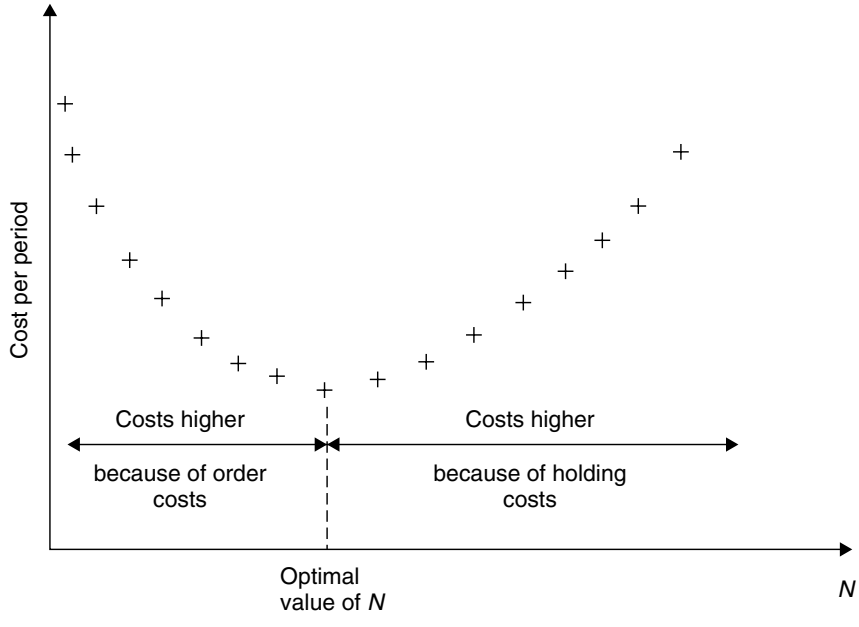


Figure 4.19 Variable cost and the number of periods in an order

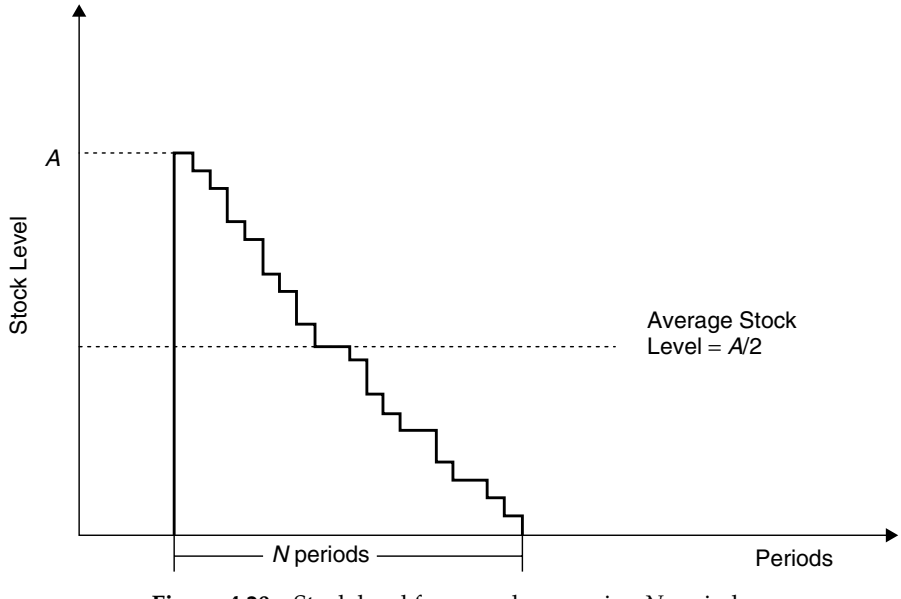


Figure 4.20 Stock level for an order covering *N* periods

The variable cost of stocking the item over the N periods is the sum of:

$$\text{reorder cost component} = RC$$

$$\text{holding cost component} = \text{average stock level } (A/2) \times \text{time held } (N) \times \text{holding cost } (HC)$$

$$= \frac{HC \times N \times \sum_{i=1}^{N} D_i}{2}$$

Adding these two components and dividing by N gives the average variable cost per period, VC_N, of:

$$VC_N = \frac{RC}{N} + \frac{HC \times \sum_{i=1}^{N} D_i}{2}$$

If you look at the short stock cycles on the left-hand side of the graph in Figure 4.19, VC_N is high because of reorder costs. Increasing N will follow the graph downward until costs reach a minimum and then start to rise. We can identify this point by comparing the cost of two consecutive values of N. If it is cheaper to order for $N + 1$ periods than for N periods, we are on the left-hand side of the graph. If we keep increasing the value of N, there comes a point where it is more expensive to order for $N + 1$ periods than N. At this point, we have found the optimal value for N and if we increase it any further, costs will continue to rise. So we can start by setting N to one, and comparing the cost of ordering for one period with the cost of ordering for two periods. If it is cheaper to order for two periods, we are on the left-hand side of the graph, so we increase the value of N to two and compare the cost of ordering for two periods with the cost of ordering for three periods. If it is cheaper to buy for three periods we are still on the left-hand side of the graph, so we increase N to three. If we continue this procedure, comparing the cost of buying for N periods with the cost of buying for $N + 1$ periods, we reach a point it becomes cheaper to buy for N periods than for $N + 1$ periods. Then we have reached the bottom of the graph and found the point of minimal cost.

We really need a simple way of comparing the variable cost of a cycle with $N + 1$ periods with the variable cost of a cycle with N periods. If we substitute $N + 1$ for N in the variable cost equation above we get:

$$VC_{N+1} = \frac{RC}{N+1} + \frac{HC \times \sum_{i=1}^{N+1} D_i}{2}$$

Now we want to find the point at which VC_{N+1} become larger than VC_N or:

$$\frac{RC}{N+1} + \frac{HC \times \sum_{i=1}^{N+1} D_i}{2} > \frac{RC}{N} + \frac{HC \times \sum_{i=1}^{N} D_i}{N}$$

After some manipulation we can simplify this to give the condition:

$$\boxed{N \times (N+1) \times D_{N+1} > \frac{2 \times RC}{HC}}$$

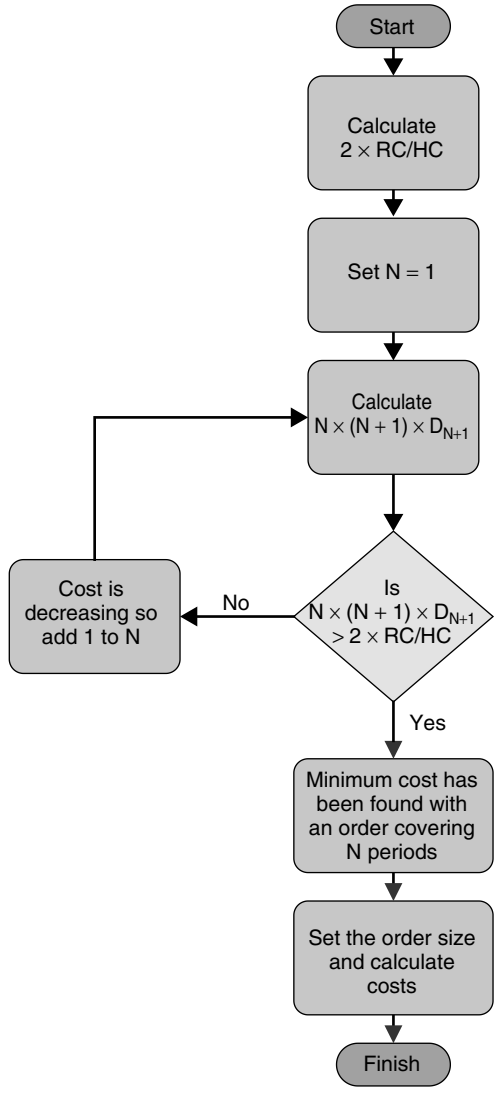


Figure 4.21 Flow chart for finding an optimal value of N

We can use this inequality in an iterative procedure to find the optimal value of N. This starts by setting N equal to 1 and checking the values in this expression. If the inequality is invalid, it is cheaper to order for two periods than for one, so we are moving down the left-hand side of the costs in Figure 4.19. Then we set N equal to 2 and check the values in the expression. If the inequality is still invalid, the costs are reducing and we are still coming down the left-hand side of the graph. We keep on increasing N, until eventually the inequality will become valid. At this point we are at the bottom of the cost curve and have found an optimal value for N. Figure 4.21 shows a flow diagram for this procedure.

Worked example

K C Low calculated her total cost of placing an order for an item and having it delivered at $1,200, with a holding cost of $50 a unit a week. She had a fixed production schedule and needed the following units over the next 11 weeks.

week	1	2	3	4	5	6	7	8	9	10	11
demand	2	4	6	9	9	6	3	4	4	6	9

Find an ordering pattern that gives a reasonable cost for the item.

Solution

We know that:

$$RC = \$1{,}200 \text{ an order}$$

$$HC = \$50 \text{ a unit a week}$$

Then following the procedure shown in Figure 4.21, we calculate:

$$2 \times RC/HC = 2 \times 120/5 = 48$$

- Taking $N = 1$, $N + 1 = 2$ and $D_2 = 4$ and calculate $N \times (N + 1) \times D_{N+1} = 1 \times 2 \times 4 = 8$. As this is less than 48 the inequality is *invalid* and we have not reached the minimum cost.
- Add one to N, giving $N = 2$, $N + 1 = 3$ and $D_3 = 6$ and calculate $N \times (N + 1) \times D_{N+1} = 2 \times 3 \times 6 = 36$. This is less than 48 so the inequality is still *invalid* and we have not reached the minimum cost.
- Add one to N, giving $N = 3$, $N + 1 = 4$ and $D_4 = 9$ and calculate $N \times (N+1) \times D_{N+1} = 3 \times 4 \times 9 = 108$. This is more than 48 so the inequality is *valid* and we have found the minimum cost with $N = 3$.

This order must arrive at the beginning of week 1 and must be big enough to cover demand in the first three weeks. This is $2 + 4 + 6 = 12$ units. Unfortunately, this result is only valid for the particular demand pattern in the first few weeks. It is not a general result and we have to repeat the procedure for every other stock cycle. So now we move on to week 4, set N equal to one, and do the calculations again for the next cycle. These calculations are easy with a spreadsheet, as illustrated in Figure 4.22. This shows that a good ordering policy is to arrange deliveries of 12 units by the start of week 1, 27 for week 4 and 14 for week 8. Adding standard information about lead time and supplier conditions would tell us exactly when to place the orders.

Inventory Control–for demand that is discrete and variable											
Notes:	Calculates the value of 2*RC/HC										
	Iteratively checks N*(N+1)*D(N+!) until this is bigger than 2*RC/HC										
Inputs											
Reorder cost	$ 1,200.00										
Holding cost	$ 50.00										
2*RC/HC	48										
Period	1	2	3	4	5	6	7	8	9	10	11
Demand	2	4	6	9	9	6	3	4	4	6	9
Calculations											
Period	1	2	3	4	5	6	7	8	9	10	11
Demand	2	4	6	9	9	6	3	4	4	6	9
N	1	2	3	1	2	3	4	1	2	3	
N*(N+1)*D(N+1)	8	36	108	18	36	36	80	8	36	108	0
Delivery	**12**			**27**				**14**			

Figure 4.22 Calculations for worked example

To find the cost of a particular policy we have to look at each stock cycle in turn and find the variable cost per period from:

$$VC_N = \frac{RC}{N} + \frac{HC \times \sum_{i=1}^{N} D_i}{2}$$

In the example above, each cycle has one order which has a cost of RC = \$1,200. The first cycle is three weeks long and the demand met is 12 units, so the average variable cost a week is:

$$VC = 1{,}200/3 + 50 \times 12/2 = \$700 \text{ a week}$$

Similarly, the next two cycles have average variable costs a week of:

- *Cycle 2*: 27 units are bought in four weeks, so VC = $1{,}200/4 + 50 \times 27/2 = \$975$ a week.
- *Cycle 3*: 14 units are bought in three weeks, so VC = $1{,}200/3 + 50 \times 14/2 = \$750$ a week.

Overall, the variable cost comes to $700 \times 3 + 975 \times 4 + 750 \times 3 = \$8{,}250$ or \$825 a week.

This method usually gives good results, but it does not guarantee optimal ones. This is largely because of the assumptions, including fixed and known costs, known demand that occurs at discrete points, and the approximations for average stock level. There is, however, one other consideration. It is sometimes possible

for the variable cost to be higher for orders covering N + 1 periods rather than N, but then fall again for N + 2 periods. To avoid this we can use the same argument as before to check that:

$$VC_{N+2} > VC_N$$

which leads to a test:

$$N \times (N+2) \times [D_{N+1} + D_{N+2}] > \frac{4 \times RC}{HC}$$

This adds a test to the procedure described above. When we find a turning point and we know it is cheaper to order for N periods than for N + 1, we simply see if it is also cheaper to order for N + 2 periods. If the inequality above is valid, we stop the process and accept the solution: if the inequality is invalid we continue the process until we find another turning point.

Summary

Organizations often have to deal with discrete, variable demand. When we know the pattern of demand in advance, we can use a simple procedure to find a good ordering policy. This procedure has to be repeated for every stock cycle.

Review questions

4.19 What do we mean by discrete, variable demand?

4.20 We defined D_i as the demand in period i. Is this correct?

4.21 Why does the rule described only give good rather than optimal solutions?

Chapter review

- This chapter described some models for inventory control. In particular, it looked at deterministic problems where all variables are known with certainty.
- The first model considered unit costs that fall in discrete steps with increasing order size. A standard procedure finds the minimum point on a valid cost curve and the corresponding order size. This standard procedure can be adjusted to deal with other situations, such as a discrete increase in reorder cost.
- Replenishment of stock often occurs at a finite rate rather than instantaneously. This gives a cycle with stock levels rising at a rate P − D during replenishment and then falling at a rate D when replenishment stops. The optimal batch size – and related results – vary from the economic order calculation by the factor $\sqrt{(P/(P-D))}$.
- Sometimes it is useful to have planned shortages, particularly when customer demand is not lost but can be met from back-orders. Back-orders inevitably

have some cost, which is time-dependent. We can use this cost to find an optimal value for the amount to be back-ordered.

- Shortages can lead to lost sales rather than back-orders. We can build a model that includes shortages and maximizes net revenue, rather than minimizing total cost. This gives a simple rule that shows whether or not to stock an item.
- There are many circumstances in which inventory items cannot be treated in isolation, such as constraints on total stock. When there is a constraint on space, we can find revised order quantities by adding an additional cost for the space used. A more formal analysis shows the effects of a constraint on the average investment in stock
- Organizations often have to deal with discrete, variable demand. We described a simple procedure to find a good ordering policy when we know the pattern of demand in advance.

Project

The project in the last chapter was to build a spreadsheet for doing some calculations with the economic order quantity. In this chapter we have illustrated some more spreadsheets for calculations with price discounts, finite replenishment rates and discrete demand. The aim of this project is to design a more sophisticated spreadsheet that can tackle a range of problems and recommend ordering policies. Decide the features that would be useful and put them into a single spreadsheet, or family of related spreadsheets. Use this to check the answers to examples in the chapter. What difficulties and problems do you meet?

In the next chapter, we will be describing some more models, so be prepared to extend your designs to deal with uncertain demand.

Problems

4.1 Demand for an item is constant at 400 units a month. The reorder cost and delivery charge amount to £1,240 and the cost of holding stock is 30 per cent of value a year. A supplier quotes the following prices

Order quantity	0–1,499	1,500–1,999	2,000–2,499	2,500 or more
Unit cost	£12.60	£12.20	£11.80	£11.20

A second supplier quotes a basic price of £12, but with a discount to £11.40 for orders of 1,500 or more. What is the best ordering policy?

4.2 Last month La Café Pigalle was sent a new price list by their merchants. The cost of their most popular wine is now:

Order quantity	0–99	100–399	400–999	1,000 or more
Price per bottle	€ 20	€ 19.40	€ 18.80	€ 18

Demand for the wine is relatively constant at 2,000 a year, delivery costs €50 and the holding cost is 40 per cent of value a year. What ordering policy would you recommend for the restaurant? If e-procurement reduces the reorder cost so that it is effectively zero, how does this affect the orders?

4.3 Demand for an item is 500 units a month, while the production rate is 1,000 units a month. Unit cost is $10, batch set-up cost is $2,000 and holding cost is $1 a unit a month. What is the optimal batch size and corresponding cost? If the production rate can be varied, how will the costs change?

4.4 Sulleman and Baring forecast demand for a component at around 2,500 units a year. They make the component internally, and it costs £500 to set up each production run with a variable cost of £30 a unit. Holding costs are 20 per cent of value a year and the production rate is 10,000 units a year. There is a lead time of two months from receiving a production requisition until finished units begin to come from the production line. Assuming that shortages are not allowed, what are the optimal batch size and reorder level for the component? What are the costs of this policy?

4.5 John van Houghton looked at an inventory item and decided that the holding cost is about 25 per cent of value a year while the shortage cost for back-orders is 150 per cent of value a year. Unit cost is €400 and reorder cost is €100. Demand is constant at 300 units a year and all shortages are met by back-orders. What is the best ordering policy for this item? What proportion of time is demand met by back-orders, and what is the cost of this policy?

4.6 Demand for an item is constant at 100 units a year. Unit cost is £50, reorder cost is £40 and holding cost is 40 per cent of value a year. Any demand that occurs when no stock remains is lost. What is the minimum selling price that makes it profitable to stock the item?

4.7 A company stocks four items with the following characteristics.

Item	Demand	Unit cost	Holding cost	Reorder cost	Space used (m^3)
1	500	200	35	80	4
2	300	100	20	100	3
3	200	300	45	120	6
4	400	200	40	80	5

Describe the effects of constrained space and investment on this system.

4.8 Andrew McGregor forecast weekly demand for MP 411 pumps as follows:

Week	1	2	3	4	5	6	7	8	9	10	11
Demand	1	6	10	21	18	7	4	3	2	4	15
Week	12	13	14	15	16	17	18	19	20	21	22
Demand	27	38	51	40	22	12	8	6	7	2	3

Each pump costs £80, it costs £800 to place an order and get delivery, and holding cost is 20 per cent of value a year. Design a reasonable ordering policy for the pump. How much does this cost a year?

Discussion questions

4.1 We have described some simple extensions to the economic order quantity, and could have looked at many more complicated models. When do you think that it is worth using a more sophisticated model, rather than using the guidelines given by a simple model?

4.2 We have looked at costs that vary in discrete steps, but how would you deal with costs that rise on a continuous sliding scale?

4.3 With a finite replenishment rate the order quantity is the economic order quantity multiplied by $\sqrt{(P/(P-D))}$. As the replenishment rate gets closer to the demand rate, this value gets larger, and the order size and average stock level both increase. But if the replenishment rate is equal to the demand we should not need any stock. What is happening?

4.4 Nobody likes waiting for a product they have decided to buy, so why would an organization deliberately work with shortages?

4.5 It is often difficult to find reliable costs for stocks. With shortages this seems almost impossible – how do you find a cost for loss of goodwill or reduced future sales? Does this mean that any analysis of shortages is inevitably based on flimsy evidence?

4.6 Every organization thinks that its problems with inventories are unique. To what extent do you think they are right?

References and further reading

Buffa, E.S. and Sarin, R.K. (1992) *Modern Production/Operations Management*, New York: John Wiley.

Chase, R.B., Aquilano, N.J. and Jacobs, R. (2000) *Production and Operations Management* (9th edition), Homewood, IL: Irwin.

Fogarty, D.W., Hoffman, T.R. and Blackstone, J.H. (1991) *Production and Inventory Management*, Dallas, TX: South-Western Publishing.

Greene, J.H. (1997) *Production and Inventory Control Handbook* (3rd edition), New York: McGraw-Hill.

Landvater, D.V. (1997) *World Class Production and Inventory Management*, New York: John Wiley.

Lewis, C.D. (1997) *Demand Forecasting and Inventory Control*, Cambridge: Woodhead Publishing.

Seetharama, L.N., McLeavy, D.W. and Billington, P. (1995) *Production Planning and Inventory Control*, Boston: Allyn and Bacon.

Silver, E.A., Pike, D.E. and Peterson, R. (1998) *Decision Systems for Inventory Management and Production Planning* (3rd edition), New York: John Wiley.

Smith, S.B. (1994) *Computer-based Production and Inventory Control*, Englewood Cliffs, NJ: Prentice-Hall.
Tersine, R.J. (1994) *Principles of Inventory and Materials Management* (3rd edition), Englewood Cliffs, NJ: Prentice Hall.
Viale, J.D. (1997) *Inventory Management*, Menlo Park, CA: Crisp Publications.
Waters, C.D.J. (1998) *A Practical Introduction to Management Science* (2nd edition), Harlow: Pearson Education.
Waters, C.D.J. (2002) *Introduction to Operations Management*, Harlow: Pearson Education.

5

Models for Uncertain Demand

Aims of the chapter

The last chapter developed some models for controlling stock when the demand and costs all took fixed values that were known in advance. In this chapter we introduce uncertainty and develop some models where variables are not known exactly but follow known probability distributions. In particular, we focus on variable demand. Many models have been developed in this area, so we will concentrate on the most widely used.

After reading this chapter you should be able to do the following:

- appreciate the types of uncertainty that come with stocks and the need for probabilistic models;
- use a marginal analysis to find the best order quantity for a single period and extend this to the newsboy problem;
- calculate an optimal stock level for intermittent demand;
- calculate order quantities with shortages;
- understand the principles of service level and safety stock;
- find the best policy for stock with Normally distributed demand;
- find an optimal policy when lead time is uncertain;
- solve problems when both demand and lead time are Normally distributed;
- describe periodic review methods and calculate target stock levels.

This chapter emphasizes:

- *uncertainty* in inventory management;
- *variables* that take uncertain values;
- *probabilistic models* for inventory control.

Uncertainty in stocks

Areas with uncertainty

The models we developed in the last chapter assumed that costs, demand, lead time and all other variables are known exactly. In other words, there is no uncertainty about the stocks. In practice, there is almost always some *uncertainty* in stocks – as prices rise with inflation, operations change, new products become available, supply chains are disrupted, competition alters, new laws are introduced, the economy varies, customers and suppliers move, and so on. From an organization's point of view, the main uncertainty is likely to be in customer demand, which might appear to fluctuate randomly or follow some long-term trend. In this chapter we will develop some models to deal with this uncertainty.

We should start by saying what we mean by 'uncertainty'. For inventory management this means that a value is not known exactly, but follows a known probability distribution. Long-term demand for a product might, for example, be Normally distributed with a mean of 10 units a week and variance of 2 units a week. We cannot say exactly what demand in a particular week will be, but know that it is a figure drawn from this distribution. Now we can classify problems according to variables that are:

- *unknown* – in which case we have complete ignorance of the situation and any analysis is difficult;
- *known* (and either constant or variable) – in which case we know the values taken by parameters and can use deterministic models;
- *uncertain* – in which case we have probability distributions for the variables and can use probabilistic or stochastic models.

You can find uncertainty in many aspects of stock. Sometimes this is due to internal operations. No matter how good the operations are, there is always some variation that can lead to uncertainty. Differences in materials, weather, tools, employees, moods, time, stress, and a whole range of other things combine to give apparently random variations. The traditional way of dealing with these was to set a tolerance in the specifications. Provided performance is within a specified range it is considered acceptable. A 250 g bar of chocolate might weigh between 249.9 g and 250.1 g and still be considered the right weight; a train might arrive within ten minutes of its published time and still be considered on time. Performance was only considered to be a problem when it fell outside the tolerance.

Taguchi (1986) pointed out that this approach has an inherent weakness. Suppose a bank sets the acceptable time to open a new account as between 20 and 30 minutes. If the time taken is 20, 25 or 30 minutes, the traditional view says that these are equally acceptable – the process is achieving its target so there is no need for improvement. But customers would probably not agree that taking 30 minutes is as good as taking 20 minutes. On the other hand, there might be little real difference between taking 30 minutes (which is acceptable) and 31 minutes (which is unacceptable). The answer, of course, is that there is not such

a clear cut-off. If you are aiming for a target, then the further you are away from the target, the worse your performance. Organizations should clearly aim at getting actual performance as close to the target as possible, and this means reducing variability. This has been a continuing theme for management over many years, and most organizations aim at low variation in the factors that are within their control.

However, there is still some variability that comes from external causes and the organization cannot control it. Every organization works within a context that is set by international trading conditions, the national economy, government policies, the business environment, competition, other organizations in the supply chain, suppliers operations, and so on. An organization cannot change these, but they are likely to affect several areas.

- *Demand*. Aggregate demand for an item usually comes from a number of separate customers. The organization has little real control over who buys their products, or how many they buy. Random fluctuations in the number and size of orders give a variable and uncertain overall demand.
- *Costs*. Most costs tend to drift upwards with inflation, and we cannot predict the size and timing of increases. On top of this underlying trend, are short-term variations caused by changes to operations, products, suppliers, competitors, and so on. Another point – that we mentioned in Chapter 2 – is that changing the accounting conventions can change the apparent costs.
- *Lead time*. There can be many stages between the decision to buy an item and actually having it available for use. Some variability in this chain is inevitable, especially if the item has to be made and shipped over long distances. A hurricane in the Atlantic, or earthquake in southern Asia can have surprisingly far-reaching effects on trade.
- *Deliveries*. Orders are placed for a certain number of units of a specified item, but there are times when these are not actually delivered. The most obvious problem is a simple mistake in identifying an item or sending the right number. Other problems include quality checks that reject some delivered units, and damage or loss during shipping. On the other hand, a supplier might allow some overage and send more units than requested. The deliveries ultimately depend on – and define – supplier reliability.

Overall, the key issue for probabilistic models is the *lead time demand*. It does not really matter what variations there are outside the lead time, as we can allow for them by adjusting the timing and size of the next order. Once we have placed an order, though, and are working within the lead time it is too late to make any adjustments. If demand outside the lead time is higher than expected, all that happens is that we reach the reorder level sooner than expected: but if demand inside the lead time is higher than expected, it is too late to make adjustments and there will be shortages. Our overall conclusion, then, is that uncertainty in demand and lead time is particularly important for inventory management.

Uncertain demand

Even when the demand varies, we could still use the mean value in a deterministic model. We know that costs rise slowly around the economic order quantity, so this should give a reasonable ordering policy. In practice, this is often true – but we have to be careful as the mean value can give very poor results.

Worked example

Demand for an item over the past 6 months has been 10, 80, 240, 130, 100 and 40 units respectively. The reorder cost is £50 and holding cost is £1 a unit a month, and any orders placed in one month become available in the following month. How good is an ordering policy based on average values?

Solution

Listing the values given:

$$D = 100 \text{ units a month (taking the average value)}$$
$$RC = £50 \text{ an order}$$
$$HC = £1 \text{ a unit a month}$$
$$LT = 1 \text{ month}$$

We can substitute these values to give:

Order quantity $Qo = \sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 50 \times 100/1)} = 100$ units

Cycle length $To = Qo/D = 100/100 = 1$ month

Reorder level $ROL = LT \times D = 1 \times 100 = 100$ units

The best policy is to order 100 units whenever gross stock (stock on hand + stock on order) falls to 100 units. On average, this happens once a month. At first, this seems a reasonable answer, but if you look more closely you can see the problems. Using this policy over the past 6 months would have given the following pattern (shown in Figure 5.1).

Month	1	2	3	4	5	6	7
Opening stock	0	90	110	0	0	0	60
Delivery	100	100	0	100	100	100	
Demand	10	80	240	130	100	40	
Closing stock	90	110	−140	−30	0	60	

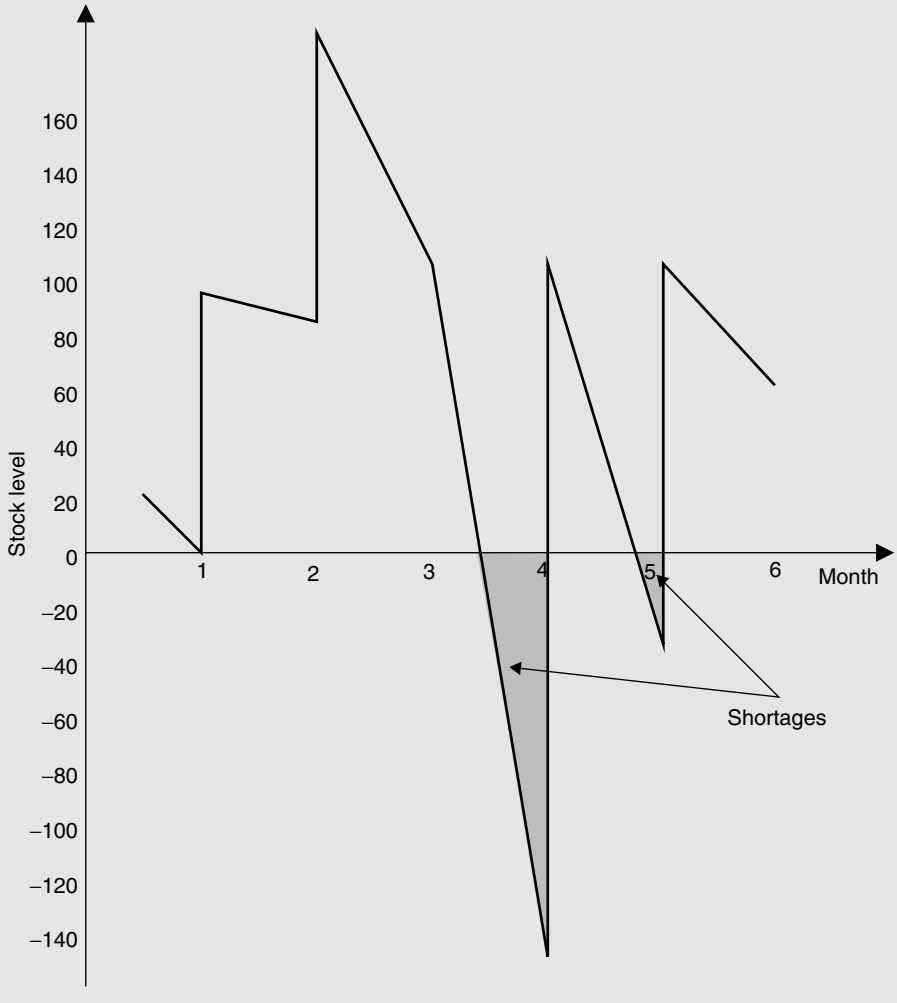


Figure 5.1 Stock level for worked example

These figures assume:

- an opening stock of zero, but an order of 100 units arriving at the start of month 1;
- orders are placed every month when the closing stock is below 100 units, and the delivery is available to meet demand in the following month;
- all unmet demand (i.e. negative closing stock) is lost.

The problem is obvious. During months 1 and 2 the stocks are unnecessarily high, but in months 3 and 4 there are shortages and lost sales. Provided

demand is below 100 units a month there are no shortage, but possible high stocks; whenever demand rises above 100 units there are shortages. In this six-month period 160 sales – or 27 per cent of the total – were lost. Not many organizations would find this acceptable.

We need to take a closer look at variable demand, and for this we will assume that the overall demand for an item is made up of small demands from a large number of customers. Then we can reasonably say that the overall demand is Normally distributed. A deterministic model will use the mean of this distribution and then calculate the reorder level as:

$$\text{reorder level} = \text{mean demand} \times \text{mean lead time}$$

Three things can happen:

- Actual demand in the lead time exactly matches expected demand. This gives the ideal pattern of stock shown in Figure 5.2(a).
- Actual demand in the lead time is less than expected demand. The resulting stock level is higher than expected, with the unused stocks shown in Figure 5.2(b).
- Actual demand in the lead time is greater than expected demand. This gives shortages, as shown in Figure 5.2(c).

With a Normally distributed demand, it is unlikely that actual lead time demand will exactly equal the expected demand (we will return to this problem of forecast errors in Chapter 7). So the actual lead time demand is likely to be either above or below the expected value. The problem is that a Normal distribution gives a demand that is higher than expected in 50 per cent of stock cycles – so we can expect shortages and unsatisfied customers in half the cycles. Very few organizations would be satisfied with this level of service, so we need to develop some models that take this uncertainty into account.

Summary

There is uncertainty in almost all inventory systems. Some of this is under the control of an organization, and this should be reduced as much as possible. More uncertainty is outside its control, including costs, demand, lead time and supplier reliability. Uncertainty in lead time demand is particularly important for inventory control.

Review questions

5.1 Where might you find uncertainty in an inventory system?

5.2 Because there is usually uncertainty in stocks, deterministic models are of little practical use. Do you think this is true?

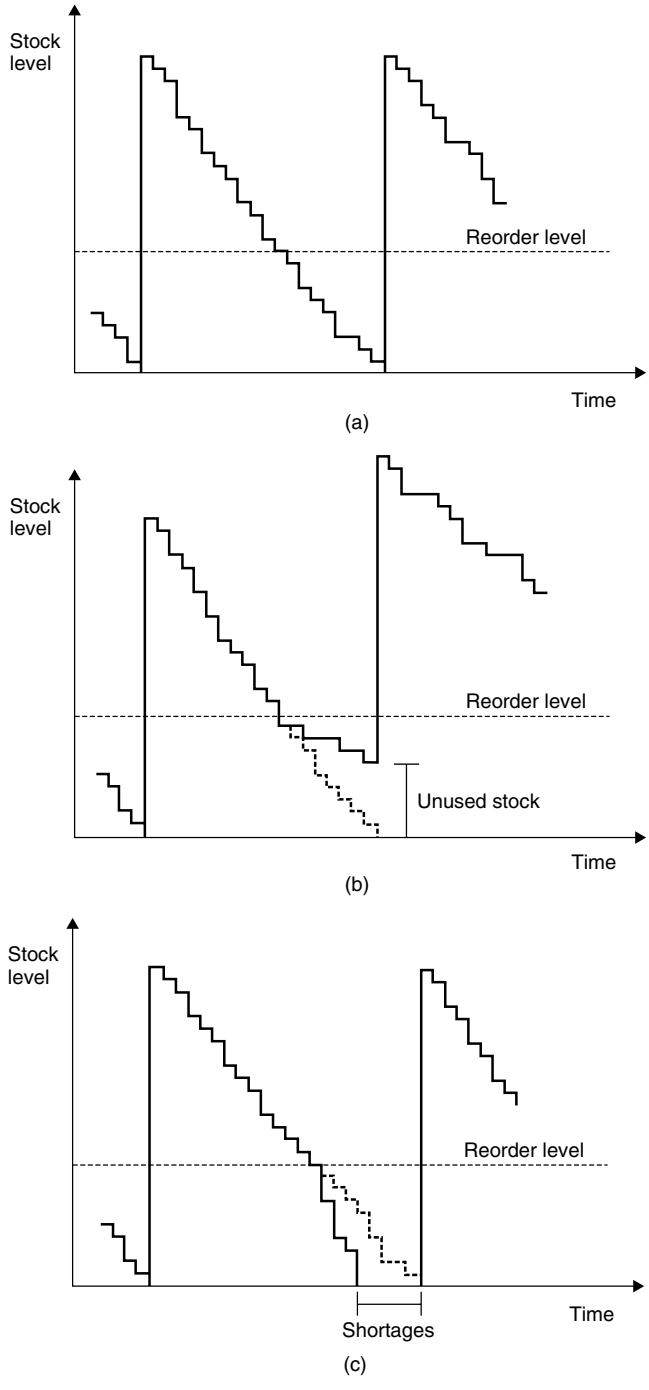


Figure 5.2 (a) The ideal case, when actual demand during lead time exactly matches expectations; (b) Unused stock when demand during lead time is less than expected; (c) Shortages when demand during lead time is greater than expected

5.3 What is the difference between variability, uncertainty and ignorance in stocks?
5.4 Why is the uncertainty in lead time demand particularly important?

Models for discrete demand

Marginal analysis

The models we have looked at so far consider stable conditions where we want the minimal cost over the long term. Sometimes, however, we need models for the shorter term and, in the extreme, for a single period. Suppose, for example, a newsagent buys a Sunday magazine from its wholesaler; it wants enough copies to meet demand on Sunday morning, but it does not want any stock afterwards. We can tackle this problem of ordering for a single cycle by using a marginal analysis, which considers the expected profit and loss on each unit.

If the demand is discrete, and we place a very small order for Q units, the probability of selling the Qth unit is high and the expected profit is greater than the expected loss. If we place a very large order, the probability of selling the Qth unit is low and the expected profit is less than the expected loss. Based on this observation, we might suggest that the best order size is the largest quantity that gives a net expected profit on the Qth unit – and, therefore, a net expected loss on the $(Q + 1)$th and all following units. Ordering less than this value of Q will lose some potential profit, while ordering more will incur net costs. So let us assume that:

- we buy a number of units, Q;
- some of these are sold in the cycle to meet demand, D;
- any units left unsold, $Q - D$, at the end of the cycle are scrapped at a lower value;
- $\text{Prob}(D > Q)$ = probability demand in the cycle is greater than Q;
- SP = selling price of a unit during the cycle;
- SV = scrap value of an unsold unit at the end of the cycle.

The profit on each unit sold is $(SP - UC)$, so the expected profit on the Qth unit is:

$$= \text{probability of selling the unit} \times \text{profit made from selling it}$$
$$= \text{Prob}(D \geq Q) \times (SP - UC)$$

And the loss on each unit scrapped is $(UC - SV)$, so the expected loss on the Qth unit is:

$$= \text{probability of not selling the unit} \times \text{loss incurred with not selling it}$$
$$= \text{Prob}(D < Q) \times (UC - SV).$$

We will only buy Q units if the expected profit is greater than the expected loss and:

$$\begin{aligned}\text{Prob}(D \geq Q) \times (SP - UC) &\geq \text{Prob}(D < Q) \times (UC - SV) \\ &\geq (1 - \text{Prob}(D \geq Q)) \times (UC - SV)\end{aligned}$$

which we can rearrange to give the general rule, that we place an order for the largest value of Q which still has:

$$\boxed{\text{Prob}(D \geq Q) \geq \frac{UC - SV}{SP - SV}}$$

Starting with a small value of Q, we can iteratively increase it, and the expected profit continues to rise while the inequality remains valid. At some point the inequality becomes invalid, showing the last units would have an expected loss, and net profit begins to fall. This identifies the best value for the order size.

Worked example

Warehouse Accessories Inc. are about to place an order for industrial heaters for a forecast spell of cold weather. They pay $1,000 for each heater, and during the cold spell sell them for $2,000 each. Demand for the heaters declines markedly after a cold spell, and any unsold units are discounted to $500. Previous experience suggests the likely demand for heaters is as follows.

Demand	1	2	3	4	5
Probability	0.2	0.3	0.3	0.1	0.1

How many heaters should the company buy?

Solution

Listing the variables that we know:

$$UC = \$1{,}000 \text{ a unit}$$
$$SP = \$2{,}000 \text{ a unit}$$
$$SV = \$500 \text{ a unit}$$

Then:

$$\frac{UC - SV}{SP - SV} = \frac{1{,}000 - 500}{2{,}000 - 500} = 0.33$$

Prob(D ≥ Q) is the cumulative probability of selling at least Q heaters, and we want the largest value of Q for which this is less than 0.33.

- If Q = 1, Prob(D ≥ 1) = 1.0. This is greater than 0.33, so the inequality is valid and we increase Q.
- If Q = 2, Prob(D ≥ 2) = 0.8. This is greater than 0.33, so the inequality is valid and we increase Q.
- If Q = 3, Prob(D ≥ 3) = 0.5. This is greater than 0.33, so the inequality is valid and we increase Q.
- If Q = 4, Prob(D ≥ 4) = 0.2. This is less than 0.33, so the inequality is no longer valid.

This identifies Q = 3 as the highest value where the inequality is valid, so the company should buy three heaters.

Newsboy problem

This marginal analysis is particularly useful for seasonal goods, and a standard example is phrased in terms of a newsboy selling papers on a street corner. The newsboy has to decide how many papers to buy from his supplier when customer demand is uncertain. If he buys too many papers, he is left with unsold stock which has no value at the end of the day: if he buys too few papers he has unsatisfied demand which could have given a higher profit. Because of this illustration, single period problems are often called *newsboy problems*. Although, it is a widely occurring problem, we will stick to the original description of a newsboy selling papers.

The marginal analysis described above is based on intuitive reasoning, but we can use a more formal approach to confirm the results. We start this by assuming the newsboy buys Q papers, and then:

- if demand, D, is greater than Q the newsboy sells all his papers and makes a profit of Q × (SP − UC) (assuming there is no penalty for lost sales);
- if demand, D, is less than Q, the newsboy only sells D papers at full price, and gets the scrap value, SV, for each of the remaining Q − D. Then his profit is D × SP + (Q − D) × SV − Q × UC.

The optimal value for Q maximizes this expected profit. To simplify the arithmetic we will assume there is no scrap value, so SV = 0, and then we get the following values for demands, profit and probabilities:

Demand	Profit	Probability
0	$0 \times SP - Q \times UC$	Prob(0)
1	$1 \times SP - Q \times UC$	Prob(1)
2	$2 \times SP - Q \times UC$	Prob(2)
: :	: : :	: :
: :	: : :	: :
: :	: : :	: :
$Q - 1$	$(Q - 1) \times SP - Q \times UC$	Prob$(Q - 1)$
Q	$Q \times (SP - UC)$	Prob(Q)
$Q + 1$	$Q \times (SP - UC)$	Prob$(Q + 1)$
: :	: : :	: :
: :	: : :	: :
: :	: : :	: :
∞	$Q \times (SP - UC)$	Prob(∞)

The total expected profit from buying Q newspapers, EP(Q), is the sum of the profits multiplied by their probabilities.

$$EP(Q) = \text{sum of}(\text{profit} \times \text{probability}) = \left(\begin{matrix}\sum(\text{expected profits} \\ \text{when } D < Q)\end{matrix} + \begin{matrix}\sum(\text{expected profits} \\ \text{when } D \geq Q)\end{matrix}\right)$$

$$= \sum_{D=0}^{Q} [D \times SP - Q \times UC] \times Prob(D) + \sum_{D=Q+1}^{\infty} Q \times [SP - UC] \times Prob(D)$$

$$= SP \times \left[\sum_{D=0}^{Q} D \times Prob(D) + Q \times \sum_{D=Q+1}^{\infty} Prob(D)\right] - Q \times UC$$

The expected profit rises with order quantity until it reaches a maximum, and then it begins to fall, as shown in Figure 5.3. So that:

if $EP(Q) - EP(Q - 1)$ is positive the profit is increasing with additional papers;

if $EP(Q) - EP(Q - 1)$ is negative the profit is decreasing with additional papers.

The best order quantity, Qo, is the point where the profit would begin to decline if another unit were bought. At this point:

$$EP(Qo) - EP(Qo - 1) > 0 \quad \text{and} \quad EP(Qo + 1) - EP(Qo) < 0$$

or $$EP(Qo) - EP(Qo - 1) > 0 > EP(Qo + 1) - EP(Qo)$$

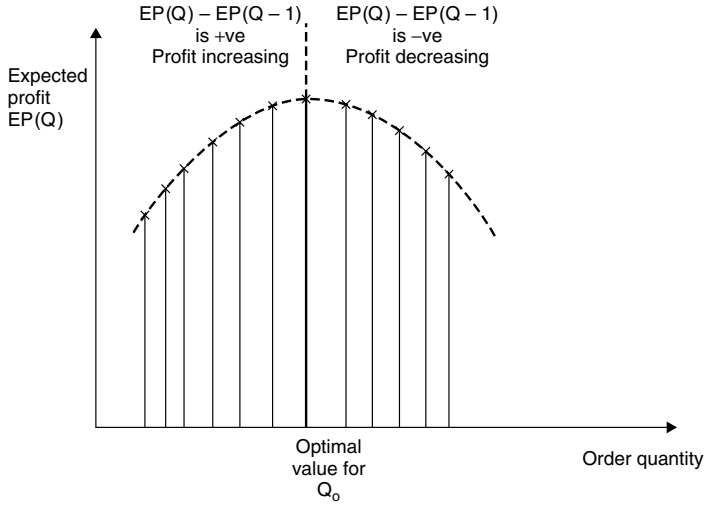


Figure 5.3 Variation of expected profit with order quantity

We have already found EP(Q), and substituting $Q-1$ for Q gives the expected profit if the newsboy buys $Q-1$ papers:

$$EP(Q-1) = SP \times \left[\sum_{D=0}^{Q-1} D \times Prob(D) + (Q-1) \times \sum_{D=Q}^{\infty} Prob(D) \right] - (Q-1) \times UC$$

With the optimal value Qo, we can substitute the values for EP(Qo) and EP(Qo − 1) and do some rearranging to find that:

$$EP(Qo) - EP(Qo-1) = SP \times \left[\sum_{D=Qo}^{\infty} Prob(D) - \frac{UC}{SP} \right]$$

Replacing Qo by (Qo + 1), we can also see that:

$$EP(Qo+1) - EP(Qo) = SP \times \left[\sum_{D=Qo+1}^{\infty} Prob(D) - \frac{UC}{SP} \right]$$

We want the point where:

$$EP(Qo) - EP(Qo-1) > 0 > EP(Qo+1) - EP(Qo)$$

or

$$SP \times \left[\sum_{D=Qo}^{\infty} Prob(D) - \frac{UC}{SP} \right] > 0 > SP \times \left[\sum_{D=Qo+1}^{\infty} Prob(D) - \frac{UC}{SP} \right]$$

Giving:

$$Prob(D \geq Qo) > UC/SP > Prob(D \geq Qo+1)$$

This is the result we found with the marginal analysis when the scrap value, SV, is zero. If we included SV in this analysis we get the final result:

$$\text{Prob}(D \geq Qo) > \frac{UC - SV}{SP - SV} > \text{Prob}(D \geq Qo + 1)$$

Worked example

In recent years the demand for a seasonal product has had the following pattern.

Units	1	2	3	4	5	6	7	8
Probability	0.05	0.1	0.15	0.2	0.2	0.15	0.1	0.05

It costs £80 to buy each unit and the selling price is £120. How many units would you buy for the season? What is the expected profit? Would your decision change if the product has a scrap value of £20?

Solution

In this example UC = £80, SP = £120 and there is no scrap value. Then UC/SP = 80/120 = 0.67. Substituting values for Q gives the following probabilities:

Q	Prob(D ≥ Q)	Prob(D)
1	1.00	0.05
2	0.95	0.10
3	0.85	0.15
4	0.70	0.20
5	0.50	0.20
6	0.30	0.15
7	0.15	0.10
8	0.05	0.05

Our rule says that we choose Qo so that:

$$\text{Prob}(D \geq Qo) > UC/SP > \text{Prob}(D \geq Qo + 1)$$

So we are looking for the two adjacent values of Q that enclose 0.67 and you can see that this happens when Q equals 4:

$$\text{Prob}(D \geq 4) > UC/SP > \text{Prob}(D \geq 5)$$

$$0.7 > 0.67 \quad > 0.5$$

The optimal policy is to buy 4 units, and then the expected profit is:

$$EP(Q) = SP \times \left[\sum_{D=0}^{Q} D \times Prob(D) + Q \times \sum_{D=Q+1}^{\infty} Prob(D) \right] - Q \times UC$$

$$EP(4) = 120 \times \left[\sum_{D=0}^{4} D \times Prob(D) + 4 \times \sum_{D=5}^{\infty} Prob(D) \right] - 4 \times 80$$

$$= 120 \times \{1.5 + 4 \times 0.5\} - 320 = £100$$

We can confirm these results using a spreadsheet, as illustrated in Figure 5.4.

If each unit has a scrap value, SV, of £20, we want:

$$Prob(D \geq Qo) > \frac{UC - SV}{SP - SV} > Prob(D \geq Qo + 1)$$

Now

$$(UC - SV)/(SP - SV) = (80 - 20)/(120 - 20) = 0.6$$

And the optimal order quantity is still 4.

	A	B	C	D	E	F	G	H
1	**Inventory Control – single period model**							
2								
3	**Notes:**	The critical ratio is N/(S+N)		where:	N = loss of not selling			
4					S = profit of selling			
5								
6	**Inputs**	Unit cost	80					
7		Selling price	120					
8								
9	**Results**	Unit profit	40		**Best order**	4		
10		Critical ratio	0.667					
11								
12	**Buy**	**Probability of selling**	**Cumulative probability**	**Expected sales**	**Expected income**	**Expected costs**	**Expected profit**	
13	0	0	1	0	0	0	0	
14	1	0.05	1	1	£ 120.00	£ 80.00	£ 40.00	
15	2	0.1	0.95	1.95	£ 234.00	£ 160.00	£ 74.00	
16	3	0.15	0.85	2.8	£ 336.00	£ 240.00	£ 96.00	
17	4	0.2	0.7	3.5	£ 420.00	£ 320.00	£ 100.00	maximum
18	5	0.2	0.5	4	£ 480.00	£ 400.00	£ 80.00	
19	6	0.15	0.3	4.3	£ 516.00	£ 480.00	£ 36.00	
20	7	0.1	0.15	4.45	£ 534.00	£ 560.00	-£ 26.00	
21	8	0.05	0.05	4.5	£ 540.00	£ 640.00	-£ 100.00	

Figure 5.4 Single period inventory model for worked example

Worked example

Zennor Package Holiday Company is about to block book hotel rooms for the coming season. The number of holidays actually booked is equally likely to be any number between 0 and 99 (for simplicity rather than reality). Each room booked costs Zennor €500 and they can sell them for €700. How many rooms should the company book if unsold rooms have no value? How many rooms should it book if unsold rooms can be sold as last-minute bookings for €200 each?

Solution

The variables given are:

$$\text{UC} = €500 \text{ a room}$$

$$\text{SP} = €700 \text{ a room}$$

Then UC/SP = 0.71.

As each number of bookings between 0 and 99 is equally likely, the probability of each number is 0.01. When unsold rooms have no value, Qo has:

$$\text{Prob}(D \geq Qo) > 0.71 > \text{Prob}(D \geq Qo + 1)$$

This occurs when Q = 29 and the optimal policy is to book 29 rooms.

When last-minute bookings allow rooms to be sold for €200, there is a 'scrap value', SV = €200. Then:

$$\frac{\text{UC} - \text{SV}}{\text{SP} - \text{SV}} = \frac{500 - 200}{700 - 200} = 0.6$$

Now we want:

$$\text{Prob}(D \geq Qo) > 0.6 > \text{Prob}(D \geq Qo + 1)$$

This happens when Qo = 40 and the optimal policy is to book 40 rooms.

Discrete demand with shortages

We can extend the newsboy problem by looking at models for discrete demand over several periods. A useful approach to this incorporates the scrap value into a general shortage cost, SC, which includes all costs incurred when customer demand is not met. We can illustrate this kind of analysis by a model that has:

- discrete demand for an item which follows a known probability distribution;
- relatively small demands and low stock levels;

- a policy of replacing a unit of the item every time one is used;
- the objective of finding the optimal number of units to stock.

You can imagine this situation with stocks of spare parts for equipment. We keep a stock of spares, and replace them whenever they are used. Now our objective is to find an optimal stock level rather than calculate an optimal order quantity. For this we need a model that balances the cost of units that are bought but never used, against the cost of shortages.

We can start by defining the following relationships.

- When an amount of stock, A, is greater than the demand, D, there is a cost for holding units that are not used. This is $(A - D) \times HC$ per unit of time.
- When demand, D, is greater than the stock, A, there is a shortage cost for demand not met. This is $(D - A) \times SC$ per unit of time.

For any particular order size, we can find the expected cost as:

$$\text{Expected cost} = \text{probability of no shortage} \times \text{holding cost for unused units} + \text{probability of a shortage} \times \text{shortage cost for unmet demand}$$

Adding this for all values of expected demand gives a total expected cost of holding A units, TEC(A):

$$TEC(A) = HC \times \sum_{D=0}^{A} Prob(D) \times (A - D) + SC \times \sum_{D=A+1}^{\infty} Prob(D) \times (D - A)$$

We can use the same reasoning as before to find an optimal stock level, Ao, with:

$$TEC(Ao) - TEC(Ao - 1) > 0 > TEC(Ao + 1) - TEC(Ao)$$

And we can follow the same approach as the newsboy problem, substituting the values for TEC(Ao), TEC(Ao − 1) and TEC(Ao + 1) and doing some manipulation to give:

$$\boxed{Prob(D \leq Ao) \geq \frac{SC}{HC + SC} \geq Prob(D \leq Ao - 1)}$$

Worked example

JP Gupta and Associates store spare parts for their manufacturing equipment. The company accountant estimates the cost of holding one unit of an item in stock for a month to be £50. When there is a shortage of the item production is disrupted with estimated costs of £1,000 a unit a month. Over the past few months there has been the following demand pattern.

Demand for the item	0	1	2	3	4	5
Proportion of months	0.8	0.1	0.05	0.03	0.015	0.005

What is the optimal stock level for the part?

Solution

We know that:

$$\text{SC} = 1{,}000 \text{ a unit a month}$$

$$\text{HC} = £50 \text{ a unit a month}$$

$$\text{SC}/(\text{HC} + \text{SC}) = 1{,}000/(50 + 1{,}000) = 0.952$$

The demand figures give the following cumulative probabilities:

Stock level, A	0	1	2	3	4	5
Prob(D = A)	0.8	0.1	0.05	0.03	0.015	0.005
Prob(D ≤ A)	0.8	0.9	0.95	0.98	0.995	1.00

We want the optimal stock level, Ao, when:

$$\text{Prob}(\text{D} \leq \text{Ao}) \geq \frac{\text{SC}}{\text{HC} + \text{SC}} \geq \text{Prob}(\text{D} \leq \text{Ao} - 1)$$

The two adjacent values of cumulative probability that enclose 0.952 have Ao equals 3:

$$\text{Prob}(\text{D} \leq 3) \geq 0.952 \geq \text{Prob}(\text{D} \leq 2)$$

$$0.98 \geq 0.952 \geq 0.95$$

This analysis relies on a value for the shortage cost SC, which can be very difficult to find. It is, though, often revealing to find the shortage cost implied by current practice. Suppose, for example, that JP Gupta and Associates in the last example actually held stocks of 4 units. We can work backwards and calculate an implied shortage cost.

The actual stock, A, is equal to 4, so we can substitute this to give:

$$\text{Prob}(\text{D} \leq \text{A}) \geq \frac{\text{SC}}{\text{HC} + \text{SC}} \geq \text{Prob}(\text{D} \leq \text{A} - 1)$$

$$\text{Prob}(\text{D} \leq 4) \geq \frac{\text{SC}}{50 + \text{SC}} \text{Prob}(\text{D} \leq 3)$$

or

$$0.995 \geq SC/(SC + 50) \geq 0.98$$

$$9{,}950 \geq SC \geq 2{,}450$$

The implied shortage cost is in the range £2,450 to £9,950 a unit a month. This is clearly a wide range, but it can be used as the basis for further discussions to agree a more reliable figure.

Intermittent demand

Many organizations have a particular problem with stocks of spare parts for equipment. These parts may be used rarely, but have such high shortage costs that they must remain in stock. Demand of this kind is called, intermittent or lumpy, with a typical pattern for consecutive periods like:

0 0 10 0 0 0 0 0 25 0 0 0 0 0 0 0 0 5 0

There are similar problems with components for batch production. The materials needed for one batch must be in stock when this batch is being worked on, but then they are not needed for other batches. The main problem is finding a reasonable forecast. One approach is to consider separately:

- expected number of periods between demands, ET;
- expected size of a demand, ED.

Then the probability of a demand in any period is 1/ET, so we can forecast demand from:

$$\text{Forecast demand} = ED/ET$$

If we know the shortage cost we can balance this against the holding cost and calculate an optimal value for A, the amount of stock that minimizes the expected total cost. Alternatively, we can look at the service level, with:

$$\begin{aligned}\text{Service level} &= 1 - \text{Prob(shortage)} \\ &= 1 - [\text{Prob(there is a demand)} \times \text{Prob(demand} > A)]\end{aligned}$$

where:

$$\text{Prob(there is a demand)} = 1/ET$$

Prob(demand > A) can be found from the distribution of demand.

As you can imagine, this kind of problem is notoriously difficult and the results are often unreliable. In practice, the best policy is often a simple rule along the lines of 'order a replacement unit whenever one is used'.

Worked example

The mean time between demands for a spare part is 5 weeks, and the mean demand size is 10 units. If the demand size is Normally distributed with standard deviation of 3 units, what stock level would give a 95 per cent service level?

Solution

We know that:

$$\text{service level} = 1 - [\text{Prob(there is a demand)} \times \text{Prob(demand} > \text{A)}]$$

and we want:

$$\text{Prob(there is a demand)} \times \text{Prob(demand} > \text{A)} = 1 - 0.95 = 0.05$$

but:

$$\text{Prob(there is a demand)} = 1/5$$

so:

$$\text{Prob(demand} > \text{A)} = 0.05/0.2 = 0.25.$$

For the Normal distribution this corresponds to 0.67 standard deviations, giving:

$$\text{A} = \text{ED} + \text{Z} \times \sigma = 10 + 0.67 \times 3 = 12 \text{ units}$$

This answer makes several assumptions, but it seems reasonable.

Summary

Many models have been developed for variable, discrete demand. Sometimes, particularly with seasonal goods, we have to control stock over a single cycle. Then we can use a marginal analysis to find the best order quantity. A more formal analysis extends this into the newsboy problem. We can use a similar approach to find the optimal stock level when there are shortage costs, such as an inventory of spare parts. Then we can extend this to look at intermittent demand.

Review questions

5.5 What is a 'single period model'?

5.6 Give some examples where single period models might be used.

5.7 What happens if a marginal analysis suggests an order size Q, but we actually buy Q + 1 units?

5.8 What values are most difficult to find in the analyses in this section?
5.9 With stocks of spare parts, why did we not find an optimal order quantity?

Order quantity with shortages

When we calculated the economic order quantity in Chapter 3, we assumed that shortages were not allowed. But we have just described a model for variable, discrete demand that included shortages. Perhaps we can combine these two and see how shortages affect the economic order quantity. An important difference in the analyses is that the EOQ assumes continuous demand, and our shortage model assumes discrete demand. For simplicity we will use discrete demand, but continuous demand only needs a small adjustment.

So we assume that demand is variable and discrete, and there is a relatively small number of shortages that are all met by back-orders. If the lead time is shorter than the stock cycle, this gives the pattern shown in Figure 5.5. The average amount of stock left at the end of a cycle depends on demand in the lead time and equals (ROL − LT × D). If we place orders of size Q, the average stock level is Q/2 plus the amount left at the end of the cycle.

$$\text{average stock level} = (\text{ROL} - \text{LT} \times \text{D}) + \text{Q}/2$$

There is a shortage of (D – ROL) whenever demand in the lead time is greater than the reorder level. Then the expected shortage per cycle is:

$$\sum_{\text{D=ROL}}^{\infty} (\text{D} - \text{ROL}) \times \text{Prob(D)}$$

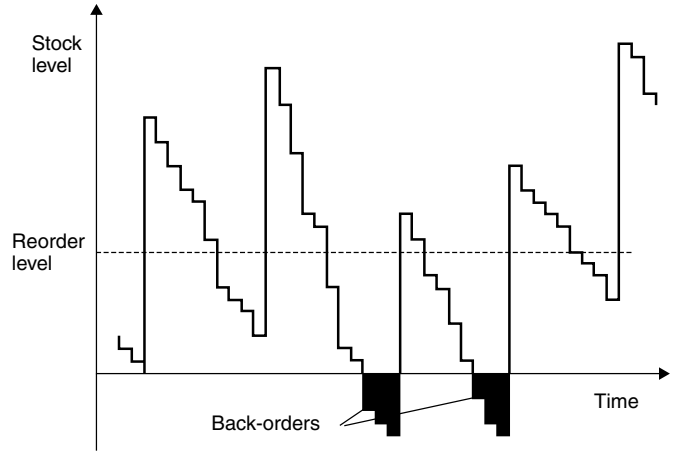


Figure 5.5 Variable demand with back-orders

Now we can find the components of variable cost per cycle, and dividing this by the cycle length gives the following costs per unit time:

- reorder cost component = number of orders (D/Q) × reorder cost (RC) = RC × D/Q;
- holding cost component = average stock (ROL − LT × D + Q/2) × holding cost (HC) = HC × (ROL − LT × D + Q/2);
- shortage cost component = expected shortage per cycle × number of cycles

$$(D/Q) \times \text{shortage cost(SC)} = \frac{SC \times D}{Q}\left[\sum_{D=ROL}^{\infty}(D - ROL) \times Prob(D)\right]$$

Adding these three components gives the expected variable cost per unit time.

$$\frac{RC \times D}{Q} + HC \times \left[ROL - LT \times D + \frac{Q}{2}\right] + \frac{SC \times D}{Q} \times \sum_{D=ROL}^{\infty}(D - ROL) \times Prob(D)$$

This equation contains two variables that are under our control, Q and ROL. Differentiating with respect to both of these and setting the derivatives to zero, gives two simultaneous equations that we can use to find optimal values. We will not do the arithmetic, but simply quote the final results as:

$$Q = \sqrt{\frac{2 \times D}{HC} \times \left[RC + SC \times \sum_{D=ROL}^{\infty}(D - ROL) \times Prob(D)\right]}$$

$$\frac{HC \times Q}{SC \times D} = \sum_{D=ROL}^{\infty} Prob(D)$$

As you can see, the optimal order quantity is greater than the economic order quantity. This is because the EOQ only balances the ordering and holding costs, but makes no allowance for shortages. Here we put more emphasis on avoiding shortages, increase the amount ordered and reduce the chance of shortages. Another point is that the equations are only valid if HC × Q is less than SC × D, or else we are looking for probabilities greater than one.

Unfortunately, the equations are not in a form that is easy to solve, so the best approach uses an iterative procedure:

1. Calculate the economic order quantity and use this as an initial estimate of Q.
2. Substitute this value for Q into the second equation and solve this to find a value for ROL.
3. Substitute this value for ROL into the first equation to give a revised value for Q.
4. Repeat steps 2 and 3 until the results converge to their optimal values.

Worked example

The demand for an item follows a Poisson distribution with mean 4 units a month. The lead time is one week, shortage cost is £200 a unit a month, reorder cost is £40, and holding cost is £4 a unit a month. Calculate optimal values for the order quantity and reorder level.

Solution

The values given, in consistent units, are:

$$D = 4 \text{ units a month}$$
$$LT = 0.25 \text{ months}$$
$$RC = £40 \text{ an order}$$
$$HC = £4 \text{ a unit a month}$$
$$SC = £200 \text{ a unit a month}$$

Following the steps of the procedure above:

Step 1. The economic order quantity is:

$$Qo = \sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 40 \times 4 \times 4)} = 8.944 \text{ units}$$

Step 2. Substituting this value into the second equation above gives:

$$\frac{HC \times Q}{SC \times D} = \sum_{D=ROL}^{\infty} Prob(D) = \frac{4 \times 8.944}{200 \times 4} = 0.045$$

We want the probability that demand is greater then the reorder level to be 0.045. As demand follows a Poisson distribution with mean 4, we can find the cumulative probabilities as follows:

Value	Probability	Prob(D ≥ Value)
0	0.018	0.999
1	0.073	0.981
2	0.147	0.908
3	0.195	0.761
4	0.195	0.566
5	0.156	0.371

Value	Probability	Prob(D ≥ Value)
6	0.104	0.215
7	0.060	0.111
8	0.030	0.051
9	0.013	0.021
10	0.005	0.008
11	0.002	0.003
12	0.001	0.001

As the Poisson distribution is discrete, we want the cumulative probability nearest to 0.045. This is 0.051, which corresponds to a value of 8, which is our initial estimate for the reorder level.

Step 3. Taking this initial reorder level and substituting it into the first equation above gives a revised value for Q;

$$Q = \sqrt{\frac{2 \times D}{HC} \times \left[RC + SC \times \sum_{D=ROL}^{\infty} (D - ROL) \times Prob(D) \right]}$$

The summation here is calculated in the following table:

D	Prob(D)	D − ROL	(D − ROL) × Prob(D)
8	0.030	0	0
9	0.013	1	0.013
10	0.005	2	0.010
11	0.002	3	0.006
12	0.001	4	0.004
Total			0.033

Then substitution gives:

$$Q = \sqrt{[(2 \times 4/4) \times (40 + 200 \times 0.033)]} = 9.654$$

Step 4. This repeats steps 2 and 3 until we find convergent values. If we substitute the new value for Q into the second equation we find a new value for ROL of:

$$(HC \times Q)/(SC \times D) = (4 \times 9.654)/(200 \times 4) = 0.048$$

This again corresponds to a reorder level of 8, so we should accept this as the final solution, with:

$$\text{reorder quantity, } Q = 9.654 \text{ (rounded to 10)};$$

$$\text{reorder level, } ROL = 8.$$

Substituting these values into the equation for variable cost per unit time shows that:

$$\begin{aligned} VC &= \frac{RC \times D}{Q} + HC \times \left[ROL - LT \times D + \frac{Q}{2} \right] \\ &\quad + \frac{SC \times D}{Q} \times \sum_{D=ROL}^{\infty} (D - ROL) \times Prob(D) \\ &= (40 \times 4)/10 + 4 \times (8 - 0.25 \times 4 + 10/2) + (200 \times 4 \times 0.033)/10 \\ &= £66.64 \text{ a month} \end{aligned}$$

In practice, the cost reductions from jointly calculating Q and ROL rather than using separate calculations are usually small. In the worked example above, simply rounding the economic order quantity to 9 units would only raise the variable cost to £66.71 a month. Sometimes, however, the calculations can make a difference, particularly when there is wide variation in lead time demand, or when the shortage cost is relatively low.

Summary

The economic order quantity does not allow shortages, so it tends to underestimate the optimal order quantity. Adding a shortage costs for discrete demand allows gives a joint calculation of order quantity and reorder level. This can sometimes give more reliable results.

Review questions

5.10 All other things being equal, does adding a shortage cost raise or lower the best order size?

5.11 When is it better jointly to calculate reorder quantities and reorder levels?

Service level

Introduction

Shortage costs are very difficult to find, but most people agree that they are high in relation to holding costs. This is why organizations are willing to incur

the relatively lower costs of holding stock, in order to avoid the higher cost of shortages. But we can go further than this and say that to avoid shortage costs, organizations should hold additional stocks – above their perceived needs – to add a margin of safety. In other words, they should hold an extra reserve of stock, knowing that it will not normally be used, but it is available when deliveries are late or demand is higher than expected. This reserve stock forms the *safety stock* (see Figure 5.6).

Earlier in the chapter we saw how, without safety stock, results based on the mean lead time demand would give shortages in 50 per cent of cycles. Now we can use a safety stock to avoid these shortages. Obviously, the larger the safety stock, the higher the chance of covering demand in a cycle and the lower the probability of a shortage. So the obvious question is, 'How much safety stock should we hold?' In principle, we can calculate the cost of shortages and balance this with the cost of holding stock. But shortage costs are difficult to find and are often little more than informed guesses. They can be so unreliable and misleading that we should view with caution any analysis using them.

An alternative approach relies more directly on managers' judgement and allows them to specify a *service level*. This is a target for the proportion of demand that is met directly from stock – or alternatively, the maximum acceptable probability that a demand cannot be met from stock. Typically an organization will specify a service level of 95 per cent, suggesting that it will meet 95 per cent of demand from stock, but will not meet the remaining 5 per cent of demand.

In practice, there are several ways of measuring the service level, with common options including:

- percentage of orders completely satisfied from stock;
- percentage of units demanded that are delivered from stock;
- percentage of units demanded that are delivered on time;

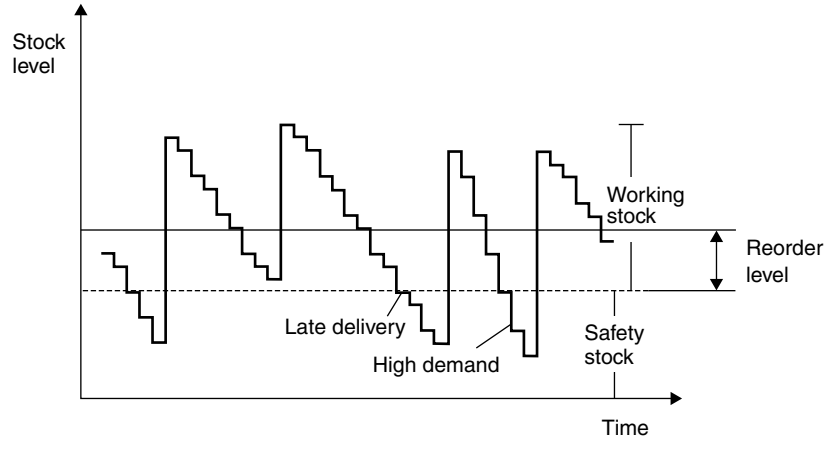


Figure 5.6 Safety stock adds a margin of security

- percentage of time there is stock available;
- percentage of stock cycles without shortages;
- percentage of item-months there is stock available.

The percentage of units demanded that are met from stock is the most obvious measure of service level. This is particularly useful in organizations like retail shops where each customer generally demands a single unit, so it also represents the percentage of satisfied customers. Unfortunately, it has the disadvantage of not taking into account the frequency of stock-outs, so a 90 per cent service level could mean a shortage once a day or once a year. In this chapter we will use *cycle service level*, which is the probability of meeting all demand in a stock cycle. Then a service level of 90 per cent means that there is a probability of 0.1 of running out of stock in a cycle, and probability of 0.9 of not running out. Again, the main problem with this definition is that it does not take into account the frequency of ordering. A 90 per cent cycle service level means there is a shortage once in 10 days for an item that is ordered daily, but only once in 10 years for an item that is ordered annually.

The way to get a higher service level is to hold more safety stock. But the critical factor in setting the amount of safety stock is the variation of lead time demand. Remember that we can allow for any variation outside the lead time, with the timing and size of the next order, but once we are within the lead time it is too late to make any changes. So the service achieved depends on the variation of lead time demand. If there is little variation, we only need low safety stocks: if there is wide variation, we need very high safety stocks to get a high service level. In principle, widely varying demand would need an infinite safety stock to ensure a service level of 100 per cent, but getting anywhere close to this can become prohibitively expensive. An organization will typically settle for a figure around 95 per cent. Often, they set different levels that reflect an item's importance; very important items may have levels around 98 per cent, while less important ones are around 85 per cent. You must, remember, though, that the choice of service level is a positive decision made by managers. They must assess all the information available and choose appropriate levels.

Worked example

In the past 50 stock cycles demand in the lead time for an item has been as follows.

Demand	10	20	30	40	50	60	70	80
Frequency	1	5	10	14	9	6	4	1

What reorder level would give a service level of 95 per cent?

Solution

Essentially, we are asking for the probability that demand in a cycle is above a specified level, so we want the cumulative distribution of lead time demand.

Lead time demand	Frequency	Probability	Cumulative probability
10	1	0.02	0.02
20	5	0.10	0.12
30	10	0.20	0.32
40	14	0.28	0.60
50	9	0.18	0.78
60	6	0.12	0.90
70	4	0.08	0.98
80	1	0.02	1.00

To achieve a service level of 95 per cent, the lead time demand must be less than the reorder level in 95 per cent of cycles. From the information given, we could set the reorder level at 70 units and get a service level of 98 per cent. Alternatively we could interpolate a figure, saying that 0.95 comes 5/8 of the way between 0.90 and 0.98, so we want a safety stock that is 5/8 of the way between 60 and 70, which is 66.25.

Summary

To avoid high shortage costs, organizations hold more stock than they expect to use. This safety stock is used when deliveries are delayed or lead time demand is higher than usual. It ensures a higher service level, which we define as the probability there are no shortages in a stock cycle. Higher safety stocks are needed for higher service levels or more variability in the lead time demand.

Review questions

5.12 What is safety stock and when is it used?

5.13 What is meant by 'service level'?

5.14 How is the service level improved?

5.15 Why is it usually impossible to guarantee a service level of 100 per cent?

Uncertain lead time demand

Uncertain demand

With deterministic models we defined the reorder level as being equal to lead time demand:

$$ROL = LT \times D$$

If the aggregate demand for an item is made up of a large number of small demands from individual customers, it is reasonable to assume the resulting demand is continuous and Normally distributed. Then, even if the lead time is constant, the lead time demand is Normally distributed and greater than the mean in half of cycles (as shown in Figure 5.7).

To get a cycle-service level above 50 per cent we have to add safety stock. We can find the size of this safety stock using the following argument. Consider an item where demand is Normally distributed with a mean of D per unit time, a standard deviation of σ, and a constant lead time of LT. Then:

- during 1 period the demand has mean D and variance of σ^2;
- during 2 periods demand has mean $2 \times D$ and variance $2 \times \sigma^2$;
- during 3 periods demand has mean $3 \times D$ and variance $3 \times \sigma^2$, and;
- during LT periods demand has mean $LT \times D$ and variance $LT \times \sigma^2$.

The mean lead time demand is $LT \times D$, variance of lead time demand is $\sigma^2 \times LT$ and standard deviation is $\sigma \times \sqrt{LT}$. The service level gives the probability that lead time demand is below the reorder level, so we can use the Normal distribution to give:

$$\text{safety stock} = Z \times \text{standard deviation of lead time}$$
$$= Z \times \sigma \times \sqrt{LT}$$

Here Z is the number of standard deviations from the mean that correspond to the specified service level. A 95 per cent service level, for example, has a probability

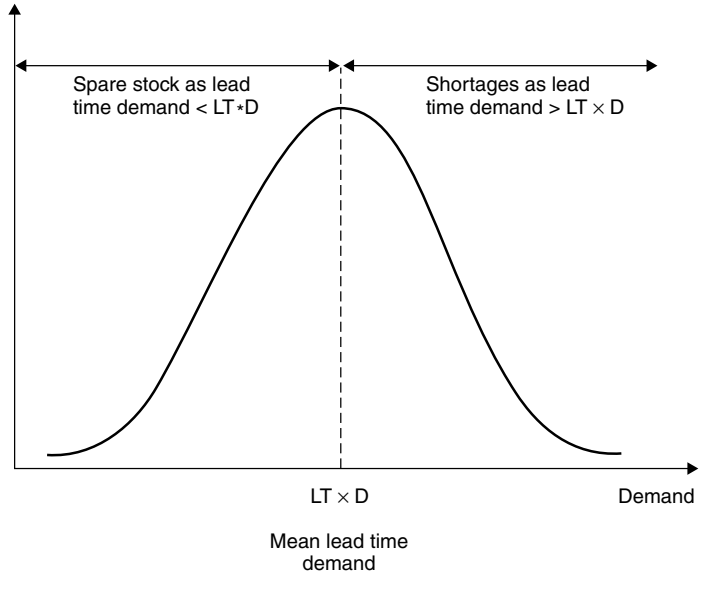


Figure 5.7 The need for safety stock with Normally distributed demand

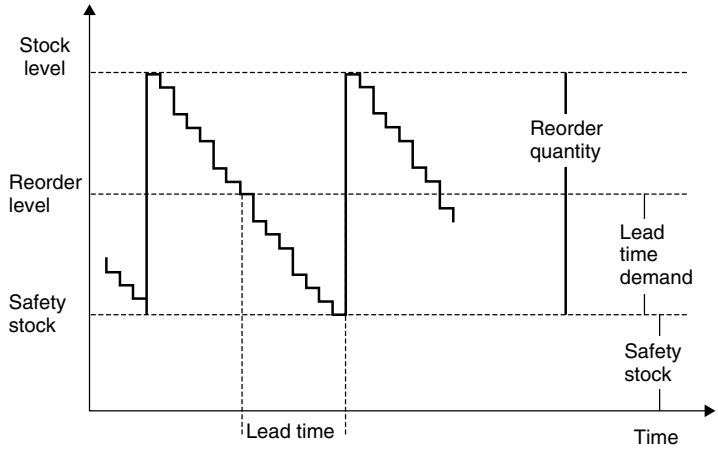


Figure 5.8 Safety stock with uncertain demand

of 0.05 that lead time demand is higher than safety stock. You can use a statistics package or Normal tables to see that a probability of 0.05 corresponds to $Z = 1.65$. Then we can calculate the safety stock, SS, from:

$$\text{safety stock} = SS = 1.65 \times \sigma \times \sqrt{LT}$$

The effect of this safety stock is to raise the reorder level (as shown in Figure 5.8) so that:

$$\begin{aligned}\text{reorder level} &= \text{lead time demand} + \text{safety stock}\\ &= LT \times D + Z \times \sigma \times \sqrt{LT}\end{aligned}$$

The following table shows how higher values of Z give higher safety stocks and lower probabilities of shortages (see also Figure 5.9).

Z	Percentage of cycles with shortages	Cycle service level (%)
0.00	50.0	50
0.84	20.0	80
1.00	15.9	84.1
1.04	15.0	85
1.28	10.0	90
1.48	7.0	93
1.64	5.0	95
1.88	3.0	97
2.00	2.3	97.7
2.33	1.0	99
2.58	0.5	99.5
3.00	0.1	99.9

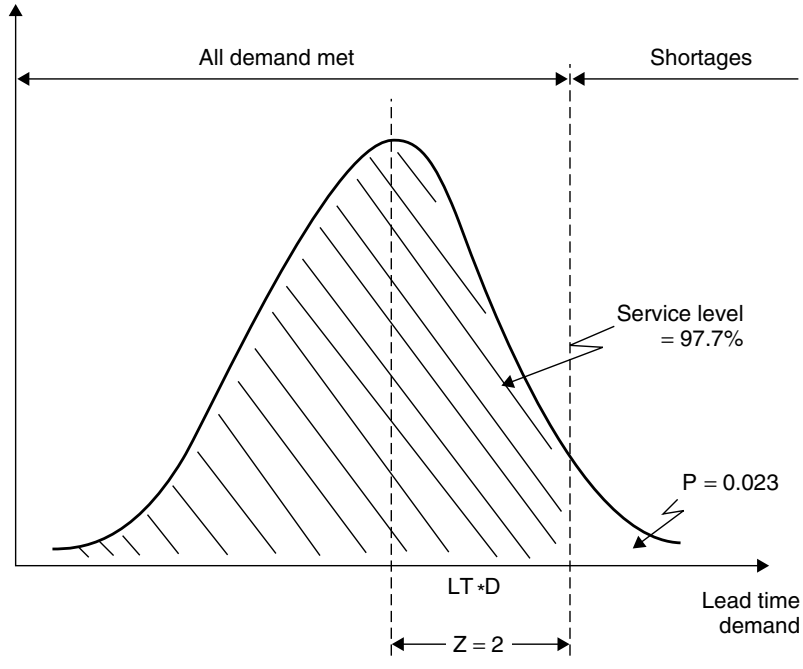


Figure 5.9 Typical service level calculation with Normally distributed demand

Worked example

A retailer guarantees a 95 per cent service level for all stock items. Stock is delivered from a wholesaler who has a fixed lead time of 4 weeks. What reorder level should the retailer use for an item that has Normally distributed demand with mean 100 units a week and standard deviation of 10 units? What is the reorder level with a 98 per cent service level?

Solution

We can find the reorder level from:

$$\text{ROL} = \text{lead time demand} + \text{safety stock} = \text{LT} \times \text{D} + \text{SS} = 4 \times 100 + \text{SS}$$

For a service level of 95 per cent, the number of standard deviations from the mean, Z, is equal to 1.64. Then:

$$\text{safety stock} = \text{Z} \times \sigma \times \sqrt{\text{LT}} = 1.64 \times 10 \times \sqrt{4} = 33 \text{ (when rounded)}$$

The reorder level becomes:

$$\text{ROL} = 400 + \text{SS} = 400 + 33 = 433 \text{ units.}$$

If the service level is increased to 98 per cent, Z = 2.05 and:

$$\text{safety stock} = 2.05 \times 10 \times \sqrt{4} = 41 \text{ units}$$

and

$$\text{reorder level} = 400 + 41 = 441 \text{ units}$$

Worked example

Polymorph Promotions plc find that demand for an item is Normally distributed with a mean of 2,000 units a year and standard deviation of 400 units. Unit cost is €100, reorder cost is €200, holding cost is 20 per cent of value a year and lead time is fixed at 3 weeks. Describe an ordering policy that gives a 95 per cent service level. What is the cost of the safety stock?

Solution

Listing the values we know:

$$D = 2{,}000 \text{ a year}$$

$$\sigma = 400$$

$$UC = €100 \text{ a unit}$$

$$RC = €200 \text{ an order}$$

$$HC = 0.2 \text{ of value held a year} = €20 \text{ a unit a year}$$

$$LT = 3 \text{ weeks}$$

Substituting these gives:

- reorder size, $Qo = \sqrt{(2 \times RC \times D/HC)} = \sqrt{2 \times 200 \times 2{,}000/20} = 200$ units
- reorder level, $ROL = LT \times D + \text{safety stock} = 3 \times 2{,}000/52 + SS = 115 + SS$
- For a 95 per cent service level $Z = 1.64$ standard deviations from the mean, so:

$$\text{safety stock} = Z \times \sigma \times \sqrt{LT} = 1.64 \times 400 \times \sqrt{(3/52)} = 158.$$

The ordering policy is to order 200 units whenever stock declines to (115 + 158 =) 273 units. Orders should arrive, on average, when there are 158 units left. The expected cost of the safety stock is:

$$= \text{safety stock} \times \text{holding cost} = 158 \times 20 = €3{,}160 \text{ a year}$$

Of course, we need not do such calculations by hand, and Figure 5.10 shows these results in a spreadsheet.

	A	B	C	D	E
1	**Inventory Control – with variable demand**				
2					
3	**Notes:**	Demand is Normally distributed with known mean and			
4		standard deviation.			
5		The safety stock is Z*sigma*sqrt(LT)			
6					
7					
8	**Inputs**	Mean demand		2,000.00	
9		Standard deviation		400	
10		Unit cost		€ 100.00	
11		Holding cost amount		€ -	
12		Holding cost percentage		20%	
13		Reorder cost		€ 200.00	
14		Lead time		3/52	
15		Service level		95%	
16					
17	**Results**	Holding cost		€ 20.00	
18	**Orders**	Order size		200.00	
19		Lead time demand		115.38	
20		Safety stock		158.03	
21		Reorder level		273.42	
22		Time between orders		0.10	
23		Number of orders a period		10.00	
24	**Investment**	Average stock		258.03	
25		Average investment		€ 25,803.24	
26		Minimum stock		158.03	
27		Minimum investment		€ 15,803.24	
28		Maximum stock		358.03	
29		Maximum investment		€ 35,803.24	
30	**Costs**	Fixed unit costs		€ 200,000.00	
31		Variable cost		€ 4,000.00	
32		Cost of safety stock		€ 3,160.65	
33		Total cost a period		€ 207,160.65	

Figure 5.10 Calculations with a Normally distributed demand

In Chapter 3 we mentioned a two-bin system for controlling some stocks. In this, one bin holds the reorder level and a second bin holds the remainder of the stock. Stock is withdrawn from the second bin until it is empty, at which point the reorder level has been reached and it is time to place an order. When demand is more variable, we can extend this to a three-bin system. Then the first bin holds the safety stock, the second holds the reorder level, and the third bin holds the rest of the stock. Stock is taken from the third bin until it is empty. At this point it is time to place an order and start using the second bin. When the second bin is empty, only the safety stock is left, and it may be time to start expediting or emergency orders.

Uncertain lead time

In the last section we described a model for uncertain demand, but constant lead time. Here we are going to change these around, and look at problems where demand is constant, but lead time is uncertain. Now we are assuming that our customers are reliable, but our suppliers are less so. Of course, suppliers try to give a good, reliable service, but even the best are affected by circumstances outside their control. If we do not allow for this uncertainty in supply and only consider the mean lead time, there are three possible outcomes.

- lead time is the expected length, and we get the ideal situation;
- lead time is shorter than expected, and there is unused stock when the delivery arrives early;
- lead time is longer than expected, and there are shortages when stock runs out before the late delivery arrives.

These stock patterns are very similar to those for variable demand, shown in Figure 5.2. In particular, we can assume that the lead time is Normally distributed and expect shortages in 50 per cent of cycles. To avoid this, we again add some safety stock and raise the reorder level. Then the probability of a shortage is simply the probability that the lead time demand is greater than the reorder level.

$$\begin{aligned}\text{Service level} &= \text{Prob}(\text{LT} \times \text{D} < \text{ROL}) \\ &= \text{Prob}(\text{LT} < \text{ROL/D})\end{aligned}$$

Worked example

Lead time for a product is Normally distributed with mean 8 weeks and standard deviation 2 weeks. If demand is constant at 100 units a week, what ordering policy gives a 95 per cent cycle service level?

Solution

We want a 95 per cent service level, so:

$$\text{Prob}(\text{LT} < \text{ROL/D}) = 0.95$$

For a Normal distribution, a probability of 0.95 corresponds to $Z = 1.64$ standard deviations. On 95 per cent of occasions lead time is less than:

$$\text{mean} + Z \times \sigma = 8 + 1.64 \times 2 = 11.3 \text{ weeks.}$$

Then the required reorder level is:

$$\text{ROL} = \text{LT} \times \text{D} = 11.3 \times 100 = 1{,}130 \text{ units.}$$

The policy is to place an order when there are 1,130 units in stock and, on average, this will arrive when there are $(1{,}130 - 8 \times 100 =)330$ units left.

Uncertainty in both lead time and demand

Now we have looked at uncertain demand with constant lead time, and uncertain lead time with constant demand. The next stage is see what happens when both lead time and demand are uncertain. If we assume that both are Normally distributed, we can use a standard result for the calculations. This says that when the demand has mean D and standard deviation σ_D, and the lead time has mean LT and standard deviation σ_{LT}, the lead time demand has mean LT × D and standard deviation, $\sigma_{LTD} = \sqrt{\text{LT} \times \sigma_D^2 + \text{D}^2 \times \sigma_{LT}^2}$.

Worked example

Demand for a product is Normally distributed with mean 400 units a month and standard deviation 30 units a month. Lead time is Normally distributed with mean 2 months and standard deviation 0.5 months. What reorder level gives a 95 per cent cycle service level? What is the best reorder quantity if reorder cost is £400 and holding cost is £10 a unit a month?

Solution

Listing the variables given:

$$\text{D} = 400 \text{ units a month}$$
$$\sigma_D = 30 \text{ units}$$
$$\text{LT} = 2 \text{ months}$$
$$\sigma_{LT} = 0.5 \text{ months}$$

Then the lead time demand has a mean of $\text{LT} \times \text{D} = 2 \times 400 = 800$ units and a standard deviation of:

$$\sigma_{LTD} = \sqrt{\text{LT} \times \sigma_D^2 + \text{D}^2 \times \sigma_{LT}^2} = \sqrt{2 \times 30^2 + 400^2 \times 0.5^2} = 204.45 \text{ units}$$

For a 95 per cent cycle service level:

$$\begin{aligned}
\text{safety stock, SS} &= Z \times \sigma_{LTD} = 1.64 \times 204.45 = 335 \text{ units} \\
\text{reorder level, ROL} &= LT \times D + SS = 400 \times 2 + 335 = 1{,}135 \text{ units} \\
\text{order quantity, Qo} &= \sqrt{(2 \times RC \times D/HC)} = \sqrt{(2 \times 400 \times 400/10)} \\
&= 179 \text{ units.}
\end{aligned}$$

Summary

When there is uncertainty, the lead time demand becomes critical. If the demand is Normally distributed, we can calculate a safety stock from $Z \times \sigma \times \sqrt{LT}$, where Z sets the service level. This raises the reorder level to $LT \times D + Z \times \sigma \times \sqrt{LT}$. A similar analysis deals with Normally distributed lead time, and we can extend this problem with uncertainty in both lead time and demand.

Review questions

5.16 Why is the lead time demand particularly important for uncertain demand?

5.17 Should the safety stock increase or decrease with increasingly variable demand?

5.18 A company keeps three weeks average demand as a safety stock. Does this seem reasonable?

5.19 What factors should affect the amount of safety stock?

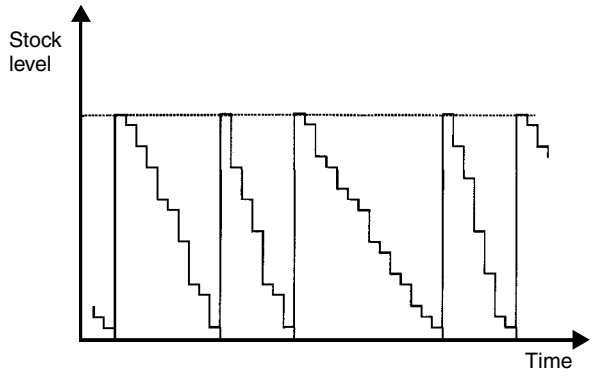
5.20 In practice, the lead time demand is always Normally distributed. Do you think this is true?

Periodic review methods

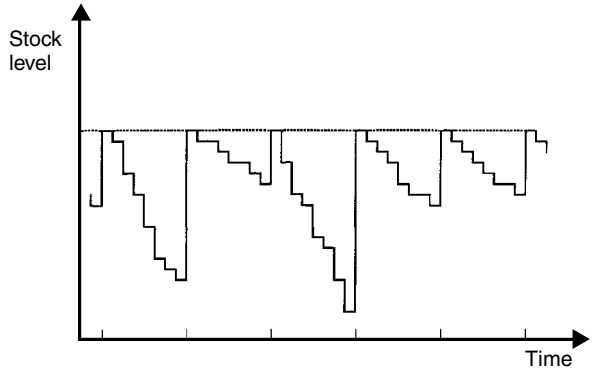
Target stock level

In Chapter 2 we mentioned two different approaches to ordering: fixed order quantity methods, where we place an order of fixed size whenever stock falls to a certain level; and periodic review methods, where we order a varying amount at regular intervals. Fixed order quantity methods allow for uncertainty by placing orders of fixed size at varying time intervals: periodic review methods allow for uncertainty by placing orders of varying size at fixed time intervals. If demand is constant these two approaches are identical, so differences only appear when the demand is uncertain (see Figure 5.11).

With a periodic review method, the stock level is examined at a specified time, and the amount needed to bring this up to a target level is ordered. A petrol



(a) Fixed order quantity



(b) Periodic review

Figure 5.11 Alternative ordering policies

station, for example, might check its stock of fuel at the end of every week, and then order enough to bring its gross stock (on hand plus on order) up to a 50,000 litres. Or a shop might examine its shelves every evening and place orders to replace items that were sold during the day. As you can see, there are two basic questions for such methods:

- How long should the interval between orders be?
- What should the target stock level be?

In practice, the order interval, T, can be any convenient period. It might be easiest to place an order at the end of every week, or every morning, or at the end of a month. If there is no obvious cycle we might aim for a certain number of orders a year or some average order size. A useful approach calculates an economic order quantity, and then finds the period that gives orders of about this size. The final decision is largely a matter for management judgement.

To find the target stock level we have to do some calculations. For convenience, we will assume that the lead time for an item is constant at LT and demand is Normally distributed. Figure 5.11(b) is an idealized view, as it shows the actual stock reaching the target stock level. In practice, there is a lead time, during which the stock falls, before the delivery arrives. The actual stock level never actually reaches the target, as shown in Figure 5.12. The size of order A is determined by the stock level at point A_1, but when this actually arrives at time A_2 stock has declined. This order has to satisfy all demand until the next order arrives at point B_2. So the target stock level has to satisfy all demand over the period A_1 to B_2, which is $T + LT$.

The demand over $T + LT$ is Normally distributed with mean of $(T + LT) \times D$, variance of $\sigma^2 \times (T + LT)$ and standard deviation of $\sigma \times \sqrt{(T + LT)}$. Using the same reasoning as before, we know that the target stock level has to be:

$$\text{Target stock level} = \text{mean demand over } (T + LT) + \text{safety stock}$$

We also know that:

$$\begin{aligned}\text{Safety stock} &= Z \times \text{standard deviation of demand over } T + LT \\ &= Z \times \sigma \times \sqrt{(T + LT)}\end{aligned}$$

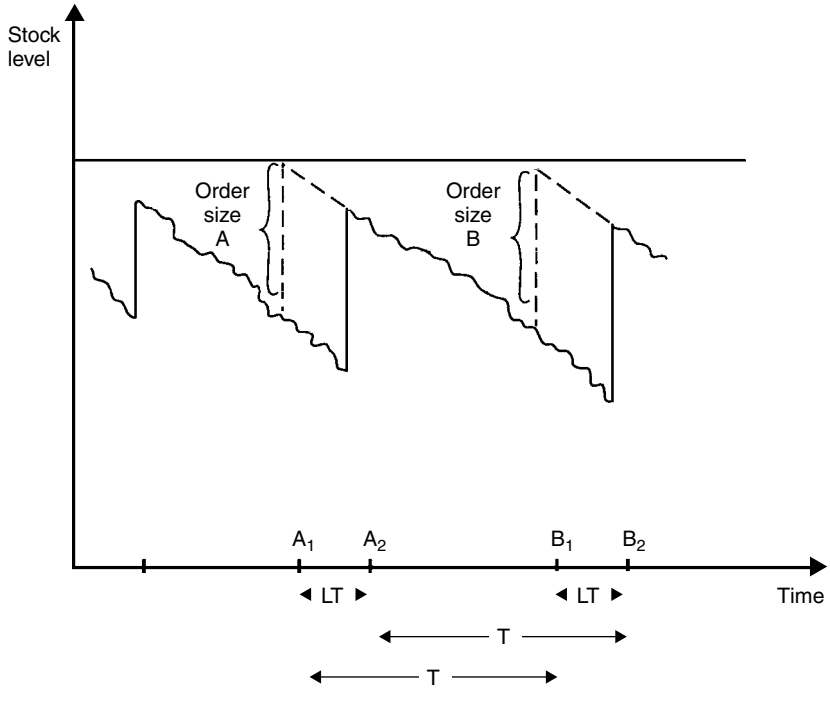


Figure 5.12 Timing of orders with periodic review

As usual, Z is the number of standard deviations from the mean corresponding to the service level. So:

$$\text{Target stock level} = D \times (T + LT) + Z \times \sigma \times \sqrt{(T + LT)}$$

This assumes that the lead time is less than the cycle length. If this is not true, the order also has to take into account the stock already on order, so that:

$$\text{Order quantity} = \text{target stock level} - \text{stock on hand} - \text{stock on order}$$

Worked example

At a recent management workshop Douglas Fairforth explained that demand for an item in his company is Normally distributed with a mean of 1,000 units a month and standard deviation of 100 units. They check stock every three months and lead time is constant at one month. They use an ordering policy that gives a 95 per cent service level, and wanted to know how much it would cost to raise this to 98 per cent if the holding cost is £20 a unit a month.

Solution

Listing the variables in consistent units for this periodic review system:

$$D = 1{,}000 \text{ units a month}$$
$$\sigma = 100 \text{ units}$$
$$HC = £20 \text{ a unit a month}$$
$$T = 3 \text{ months}$$
$$LT = 1 \text{ month}$$

For a 95 per cent safety stock Z is 1.64. Then:

$$\text{safety stock} = Z \times \sigma \times \sqrt{(T + LT)} = 1.64 \times 100 \times \sqrt{(3 + 1)} = 328 \text{ units}$$

$$\text{target stock level} = D \times (T + LT) + \text{safety stock} = 1{,}000 \times (3 + 1) + 328$$
$$= 4{,}328 \text{ units}$$

Every three months, when it is time to place an order, the company examines the stock on hand and places an order for:

$$\text{order size} = 4{,}328 - \text{stock on hand.}$$

If, for example, there were 1,200 units in stock the order would be for $4{,}328 - 1{,}200 = 3{,}128$ units. The cost of holding the safety stock $= SS \times HC = 328 \times 20 = £6{,}560$ a month

	A	B	C	D	E	F	G
1	**Inventory Control – periodic review system**						
2							
3	**Notes:**	Demand is Normally distributed with known mean and					
4		standard deviation.					
5		The target stock level is D*(T + LT) + Z*sigma*sqrt(T + LT)					
6							
7							
8	**Inputs**	Mean demand		1,000.00			
9		Standard deviation		100			
10		Unit cost		£ 100.00			
11		Holding cost amount		£ 20.00			
12		Holding cost percentage		0%			
13		Reorder cost		£ 200.00			
14		Lead time		1.00			
15		Service level		95%			
16		Period between orders		3.00			
17							
18	**Results**	Holding cost		£ 20.00		**Actual stock**	**Order**
19	**Orders**	Safety stock		328.97		250	4,078.97
20		Orders a period		0.33			
21		Demand over T + LT		4,000.00			
22		Target stock level		4,328.97			
23	**Investment**	Average stock		1,828.97			
24		Average investment		£182,897.06			
25		Minimum stock		328.97			
26		Minimum investment		£ 32,897.06			
27		Maximum stock		3,328.97			
28		Maximum investment		£332,897.06			
29	**Costs**	Fixed unit costs		£100,000.00			
30		Variable cost		£ 36,646.08			
31		Cost of safety stock		£ 6,579.41			
32		Total cost		£136,646.08			

Figure 5.13 Calculations for periodic review

If the service level is increased to 98 per cent, Z = 2.05 and:

$$\text{safety stock} = 2.05 \times 100 \times \sqrt{4} = 410$$

target stock level is then 4,410 units and the cost of the safety stock is $410 \times 20 = £8{,}200$ a month.

These calculations are easy on a spreadsheet, as shown in Figure 5.13.

Advantages of each method

We cannot say that a fixed order quantity method or a periodic review method is inevitably better than the other. Each has advantages and performs better in different circumstances. The choice may not be obvious, and is largely a matter of management preference.

The main benefit of a periodic review method is that it is simple and convenient to administer. There is a routine where stock is checked at regular times, orders are placed, delivery is arranged, goods arriving are checked, and so on. This is particularly useful for cheap items with high demand. The routine also means that the stock level is only checked at specific intervals and does not have to be monitored continuously (unlike fixed order quantity methods where stock has to be checked continuously to note the point when it falls to the reorder level). Another advantage of periodic review methods is the ease of combining orders for several items into a single order. The result can be larger orders that encourage suppliers to give price discounts.

On the other hand, a major advantage of fixed order quantity methods is that orders of constant size are easier to administer than variable ones. Suppliers know how much to send and the administration and transport can be tailored to specific needs (perhaps supplying a truck load at a time). They also mean that orders can be tailored to the needs of each item – unlike periodic review methods that commonly use the same period for many diverse items, and items with low demands are ordered as frequently as those with high demands.

Perhaps the major advantage of fixed order quantity methods is that they give lower stocks. The safety stock has to cover uncertainty in the lead time, LT, while the safety stock in a periodic review method has to cover uncertainty in the cycle length plus lead time, $T + LT$. This allows smaller safety stock and hence lower overall stocks.

Sometimes it is possible to get the benefits of both approaches by using a hybrid method. Two common types of hybrid are:

- *Periodic review with reorder level*. This is similar to the standard periodic review method, but we only place an order if stock on hand is below a specified reorder level. If the stock on hand is above the reorder level, we do not place an order this period, but wait until next period.
- *Reorder level and target stock*. This is a variation of the fixed order quantity method which is useful when individual orders are large, and might take the stock level well below the reorder level. Then, when stock falls below the reorder level, we do not order for the economic order quantity, but order an amount that will raise current stock to a target level. This is sometimes called the *min-max system* as gross stock varies between a minimum (the reorder level) and a maximum (the target stock level).

Summary

A periodic review method places orders of variable size at regular intervals. The amount ordered is enough to raise stock on hand plus stock on order to a target level, TSL, where:

$$TSL = D \times (T + LT) + Z \times \sigma \times \sqrt{(T + LT)}$$

Both fixed order quantity methods and periodic review methods have advantages in different circumstances.

Review questions

5.21 How is the order size calculated for a periodic review method?

5.22 Will the safety stock generally be higher for a fixed order quantity method or a periodic review method?

5.23 Why would a company decide to use a periodic review method?

Chapter review

- The last chapter developed some models for inventory control where the variables were known with certainty. This chapter extended the range of models by including variables with uncertain values.
- There is uncertainty in almost all stocks. Some of this is under the control of an organization, and this should be reduced as much as possible. More uncertainty is outside the organization's control, including costs, demand, lead time and supplier reliability.
- Uncertainty in lead time demand is particularly important for inventory control.
- Many models have been developed for variable, discrete demand. Sometimes, particularly with seasonal goods, we have to control stock over a single cycle. A marginal analysis can find the best order quantity, while a more formal analysis extends this to the newsboy problem.
- We can use a similar approach to find the optimal stock level when there are shortage costs, such as an inventory of spare parts. We can extend this, in turn, to look at intermittent demand.
- To avoid high shortage costs, organizations hold more stock than they expect to use. This safety stock is used when deliveries are delayed or lead time demand is higher than usual. Safety stock raises the reorder level and gives higher service levels, with more uncertainty and higher service levels needing more safety stock.
- There is often uncertainty in demand. If the overall demand is formed from many separate, small demands, it is reasonable to assume that it is continuous and Normally distributed. Then we can calculate a safety stock and corresponding reorder level.
- The lead time is often uncertain. With Normally distributed lead time, we can find a reorder level that gives a desired service level. We can extend this analysis to problems with uncertainty in both lead time and demand.
- A periodic review method places orders of variable size at regular intervals. The amount ordered is enough to raise stock on hand plus stock on order to a target level.
- Both fixed order quantity methods and periodic review methods have advantages in different circumstances.

Project

There is a lot of information about inventory control on the Web. Some useful sites are:

www.apics.org	American Production and Inventory Control Society
www.inventorymanagement.com	Centre for Inventory Management
www.cris.com	Inventory Control Forum
www.poms.org	Production and Operations Management Society
www.iomnet.org.uk	Institute of Operations Management
www.theorsociety.com	Operational Research Society

Consultants and software suppliers run many other sites. The aim of this project is to do a survey of information available on the Web. See what sites there are and what information they contain. What information can you find about the actual use of inventory control models?

Problems

5.1 In a given period demand for an item is equally likely to be any value between 100 and 199 units. Each unit of the item costs £75 and can be sold for £100. At the end of the period all unsold stock is passed to a recycler who pays £25 a unit. How many units would you buy for the period?

5.2 At the beginning of December Southern Conifers Limited employ a contractor to cut enough trees to meet the expected demand for Christmas trees. They sell these to a local wholesaler in batches of 100. Over the past few years the demand has been as follows.

Batches	1	2	3	4	5	6	7	8	9
Probability	0.1	0.1	0.2	0.2	0.1	0.1	0.1	0.1	0

If it costs £16 to cut and trim a tree that sells for £25, how many trees should the company cut down?

5.3 Over the past few months, Kepfler and Associates have collected the following figures for demand of an electric transformer that they keep as a spare part for assembly equipment.

Demand	1	2	3	4	5	6	7	8	9	10
% of months	50	20	10	5	5	3	3	2	1	1

The cost of holding one transformer in stock for a month is €60. Any shortage of the part is estimated to cost €1,000 a unit a month. What is the best stock level for the transformer? If the store currently has eight transformers in stock, what is the implied shortage cost?

5.4 The demand for an item is Normally distributed with mean of 40 units and standard deviation of 4 units a month. The lead time is one month, shortage cost is £200, reorder cost is £40, and holding cost is £4. What are optimal values for the order quantity and reorder level?

5.5 Demand for an item is Normally distributed with a mean of 400 units a week and a standard deviation of 60 units. Ordering and delivery cost $300, holding cost is $12 a unit a year and lead time is constant at 3 weeks. Describe an ordering policy that gives a 95 per cent cycle-service level. What is the cost of holding the safety stock in this case? How much would costs rise if the service level is raised to 98 per cent?

5.6 Lee Brothers advertise a 90 per cent cycle-service level for all stock items. Stock is replenished from a single supplier who guarantees a lead time of 4 weeks. What reorder level should they use for an item that has a Normally distributed demand with mean 2,000 units a week and standard deviation of 200 units? What is the reorder level for a 95 per cent cycle-service level? What is the effect of moving to a periodic review method with interval of 2 weeks?

5.7 Suppose that you have recorded the lead time for an item as follows:

Percentage of orders	5	10	15	40	15	10	5
Lead time (weeks)	1	2	3	4	5	6	7

The demand is constant at around 10 units a week and the economic order quantity has been calculated at 100 units. What ordering policy gives a cycle service level of 95 per cent for the item?

5.8 Demand for an item is 100 units a week with a standard deviation of 10 units. Lead time is one week and the reorder level used is 115 units. What is the probability of running out of stock?

5.9 Sherman (Wholesale) plc meet supply an item that has a mean demand of 200 units a week and standard deviation of 40 units. Stock is checked every 4 weeks and lead time is constant at 2 weeks. Each unit costs about £10 a week to store, what can you say about their options for inventory control?

Discussion questions

5.1 By definition, we cannot predict uncertain things. What, then, is the point of building models that contain uncertainty?

5.2 Service level models assume that we can define an acceptable level of service. But surely we should be aiming for perfect service, in the same way that Total Quality Management aims for perfect quality. Is this a major flaw in these models?

5.3 If we include shortage costs, we find that the optimal order quantity is higher than the EOQ. The reasoning is that orders are bigger to avoid shortages. But the EOQ calculation assumes that shortages are so expensive that they must never occur. Has some calculation gone wrong?

5.4 As they give lower stocks, fixed order quantity methods should be used whenever possible. Do you think that this is true?

5.5 What types of uncertainty are important in real inventory methods, but have not been included in the models we have described? How could we add these factors?

5.6 What features would you expect to see in a computerized inventory control package? Look at some commercial packages and compare the features they offer. (You can find information about many packages on the Web.)

5.7 If we kept removing the assumptions in our analyses we would end up with a model that would accurately describe the operations of any stocks. Admittedly this model might be quite complex, but an organization would simply have to substitute the appropriate values to find its best inventory policy. Is this a realistic view?

References and further reading

Buffa, E.S. and Sarin, R.K. (1992) *Modern Production/Operations Management*, New York: John Wiley.

Chase, R.B., Aquilano, N.J. and Jacobs, R. (2000) *Production and Operations Management* (9th edition), Homewood, IL: Irwin.

Fogarty, D.W., Hoffman, T.R. and Blackstone, J.H. (1991) *Production and Inventory Management*, Dallas, TX: South-Western Publishing.

Greene, J.H. (1997) *Production and Inventory Control Handbook* (3rd edition), New York: McGraw-Hill.

Landvater, D.V. (1997) *World Class Production and Inventory Management*, New York: John Wiley.

Lewis, C.D. (1997) *Demand Forecasting and Inventory Control*, Cambridge: Woodhead Publishing.

Porteus, E.L. (2002) *Foundations of Stochastic Inventory Theory*, Stanford, CA: Stanford University Press.

Seetharama, L.N., McLeavy, D.W. and Billington, P. (1995) *Production Planning and Inventory Control*, Boston: Allyn and Bacon.

Silver, E.A., Pike, D.E. and Peterson, R. (1998) *Decision Systems for Inventory Management and Production Planning* (3rd edition), New York: John Wiley.

Smith, S.B. (1994) *Computer-based Production and Inventory Control*, Englewood Cliffs, NJ: Prentice-Hall.

Taguchi, G. (1986) *Introduction to Quality Engineering*, Tokyo: Asian Productivity Association.

Tersine, R.J. (1994) *Principles of Inventory and Materials Management*, Englewood Cliffs, NJ: Prentice Hall.
Viale, J.D. (1997) *Inventory Management*, Menlo Park, CA: Crisp Publications.
Waters, C.D.J. (1998) *A Practical Introduction to Management Science* (2nd edition), Harlow: Pearson Education.
Waters, C.D.J. (2002) *Introduction to Operations Management*, Harlow: Pearson Education.

Part III

Information for Inventory Management

6

Sources of Information

Aims of the chapter

The last three chapters described some models for inventory control. To work properly, these models need a lot of accurate information about costs, demand, lead times, and so on. This is collected, analysed and presented by a management information system. In this chapter we discuss some of the information flows needed to support inventory management. In particular, we concentrate on the information provided by accounting, procurement and warehousing. The next two chapters talk about other sources of information.

After reading this chapter you should be able to do the following:

- discuss the information needed for inventory management;
- outline the design of an inventory management information system;
- consider the information supplied by accounting;
- do an ABC analysis of stock;
- describe the function of procurement and its role in inventory management;
- outline the impact of e-commerce;
- discuss the role of warehousing;
- describe important factors in the design of a warehouse.

This chapter emphasizes:

- *information* needed for inventory management;
- features of an inventory management *information system*;
- *sources* of relevant information.

Inventory management information systems

Types of information

In the last three chapters we have developed a series of models for inventory control. These were based on a few variables, so you might get the impression that

inventory managers only need to estimate a few some costs and demands, do some calculations and keep an eye on stock levels. But this is clearly a gross simplification. In reality, inventory managers need a huge amount of information – not just about the items, their specification and use, but also about business plans and strategies, suppliers, customers, promotions, storage requirements, constraints, prevailing conditions of trade, contract law, transport arrangements, competitors, business environment, changing conditions, international trade, special requirements, and so on. This information is collected from many sources, and is presented by the organization's management information system (MIS).

> The MIS controls the flow of information throughout an organization, and makes sure that everyone has the information they need to work properly.
>
> It collects, checks, organizes, stores, analyses and presents information in the most appropriate formats to everyone who needs it.

The systems take a huge variety of formats, and range from casual ones (often little more than meetings in corridors and coffee lounges) to very formal ones that give a rigid structure and procedure for making decisions. Most organizations go beyond the basic systems, and use a range of knowledge bases, decision support systems, expert systems, executive information systems, and other systems and tools to give positive assistance with decisions.

Sometimes it is convenient to refer to an *inventory management information system* which consists of all the parts of an MIS that are specifically concerned with inventory management. This collects information from both within the organization and from external sources and presents it to inventory managers. You can get an idea of the range of information needed from the following list:

- *Business environment*, which gives the overall context for stocks. Inventory managers need information about the state of competition, changing circumstances, likely shortages or supply difficulties, state of the economy, government policies and legislation, inflation, exchange rates, tariffs, duties, taxes and quotas, and a range of other external factors that might affect stocks.
- *Organizational strategies*, which are the internal policies that show what the organization is doing in the long term and the support needed from inventory management. They include details of objectives, planned operations and products, priorities, financing arrangements, trading relationship, and a range of other internal factors that might affect stocks.
- *Performance targets*, which identify the factors that are important to the organization – and consequently to inventory management – and the levels of performance to be achieved. They describe aims for customer service, productivity, stock turnover, return on inventory investment, overall investment in stock, improvements, balance between competing objectives, etc.
- *Plans for operations and production*. Stocks are needed to support operations within an organization, so inventory managers obviously need to know as much

as possible about these planned operations, including material requirements, new products, changes to existing products, promotions, size of demands, etc.

- *Costs*. Our independent demand inventory models are based on reliable figures for the four cost components of unit, reordering, holding and shortage. As well as these, managers need detailed breakdowns of other relevant costs, such as opportunity costs, financial charges, transport, storage, packaging, handling, insurance, obsolescence, internal movement, delays, distribution, etc.
- *Customer information*. Inventory managers pass materials to customers (either internal or external). They must have a clear picture of these customers, and this needs a lot of information about alliances, trading arrangements, identification of existing and new customers, contact details, salespeople used, locations, products demanded, size of orders, regular orders, preferred delivery modes, purchases from competitors, sensitivity to prices, returns, complaints, attitude towards shortages, back-orders, service level and lead times demanded, stocks held, purchasing system used, attitudes towards shared systems and integration, trading history, payment method, financial security, credit terms requested, etc.
- *Demand patterns*. To have enough stock to meet likely patterns of future demand, we need forecasts for each period, complete histories of past demand, changing patterns over time, reasons for changes, amount of variation and uncertainty, factors that affect demand, special promotions and marketing campaigns, regular orders, advance sales, etc.
- *Supplier information*, which gives all relevant information about the supply of materials, including preferred suppliers, alliances and other trading arrangements, attitude towards shared systems and integration, locations, products supplied, quality, purchasing system used, lead times, reliability, trading history, financial security, alternative suppliers, process and purchase conditions, credit terms offered, attitude towards returns, payment method, etc.
- *Product details*, which identify the exact item and include identification codes, description, design features and specifications, limitations, unit prices, discounts available, special conditions for storage or supply, quality available, weight and size, packaging, available suppliers, order restrictions, relative make/buy costs, and alternative products.
- *Warehousing information*, which shows where stocks are held, including locations, stock levels in each location, space available, facilities and conditions, material handling equipment, security, ownership, restricted access, costs and performance.
- *Transport*, including the types of transport available, types of vehicle, special requirements and facilities, access, speed of delivery, scheduled services and frequency, reliability, costs, capacities, back-hauls and reverse logistics.
- *Stock information*. This gives the details of all current stocks held. These details might include identification codes, current holdings, locations, inventory policies, reorder levels and quantities, deliveries due, suppliers, quality, costs,

units reserved for known orders, allowed variation in stock, maximum and minimum levels, age of units, obsolescent and slow moving items, insurance, special conditions, allocation of stocks to orders, back-orders, etc.

- *Orders outstanding*, which are all the orders with deliveries expected in the near future, details of materials in transit (coming from suppliers, going to customers, or moving between facilities), routes, transport operators, current locations, quantities, due dates, reliability, special handling required, potential difficulties, back-orders, urgent deliveries, etc.

This is, of course, only a partial list of the information that inventory managers need, but you can see that a large stock will need a huge amount of information. This all has to be accurate, reliable, with the right level of detail, complete and precise enough for decisions, in time, understandable, and so on.

It is easiest to imagine an inventory management information system as collecting information from different functions in an organization, storing and processing this information, passing this to inventory managers and returning reports to other functions. These reports include decisions, orders, costs, stock levels, purchases, summary statistics and any other information needed about stocks. Figure 6.1 illustrates these main information flows between inventory management and related functions.

Transaction recording

Inventory management information systems come in many different forms to suit individual circumstances, but they all have common features. They are, for example, all based on a centralized store of information that is continuously updated by new stock transactions. Every time a unit is sold or moved it affects the stock holdings, so the system has to track every detail of stock movement, from orders sent to suppliers through to payments received for finished products delivered to customers. It is difficult to give a general picture, but a typical transaction recording system might do the following:

- maintain accurate records of current stocks, customers, suppliers, etc.;
- update these records after each order, delivery, withdrawal or other transaction;
- check and validate all data used by the system.

This data validation and transaction recording stage comes before the main information processing, as shown in Figure 6.2. When you remember that an inventory of any size has tens of thousands of items and meet thousands of orders a day, you can see that even this part of the system can be big and complicated. Thankfully, the procedures are fairly routine, and there are very reliable systems that give few problems or errors.

The information collected by the transaction recording system can be presented to managers in detail, or summarized in reviews. Just as importantly, it can be

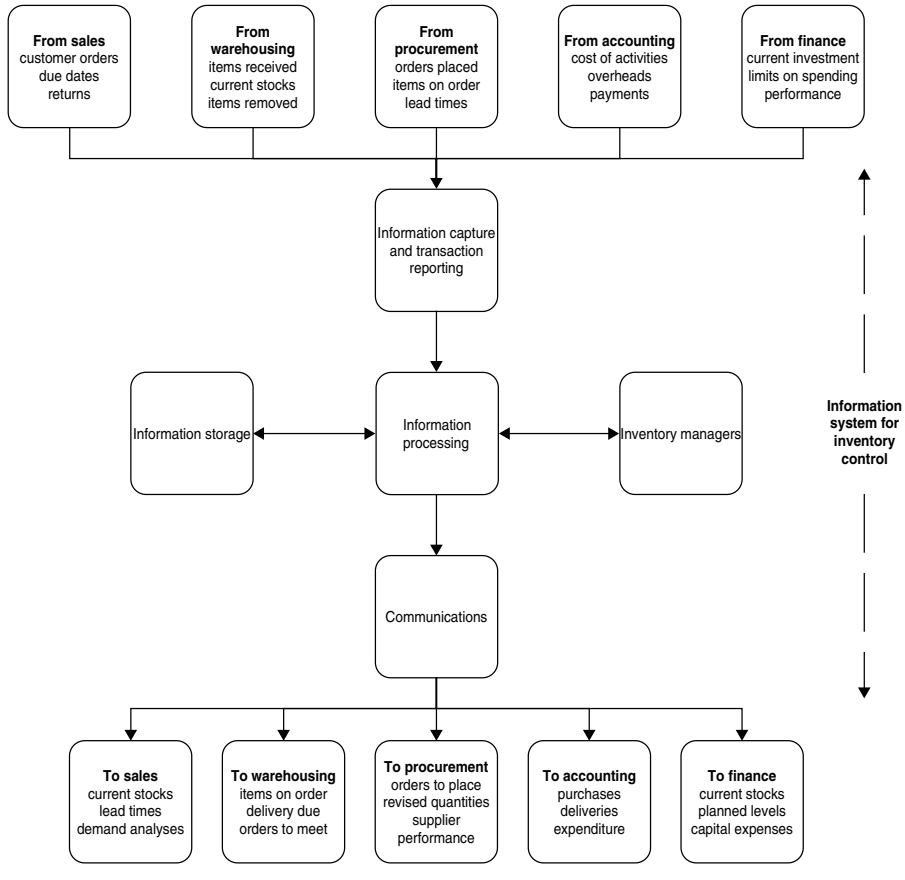


Figure 6.1 Outline of an information system for inventory management

passed to other systems for routine processing. The inventory control system, for example, can use it automatically to do the following:

- calculate variables, including forecasts of demand and lead time, reorder levels, order quantities, costs, etc.;
- prepare and transmit orders when stocks fall to reorder levels;
- monitor order progress, with follow-up and expediting of overdue orders;
- check orders, credit and payment facilities of customers;
- check deliveries, clear of invoices from suppliers and arrange payment;
- report exceptional circumstances that need management attention;
- summarize data in management reports;
- update all parameters used by the system.

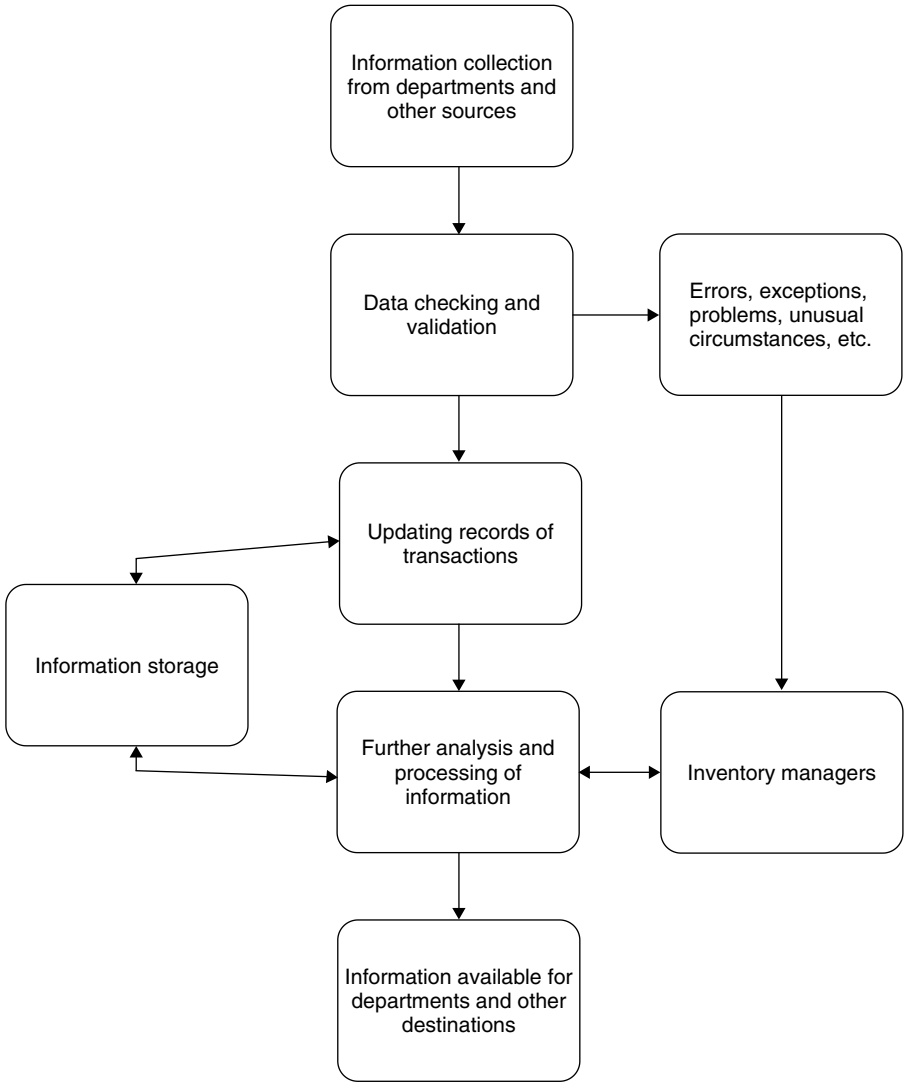

Information collection from departments and other sources
Data checking and validation
Errors, exceptions, problems, unusual circumstances, etc.
Updating records of transactions
Information storage
Further analysis and processing of information
Inventory managers
Information available for departments and other destinations

Figure 6.2 Some elements in transaction recording

One important point is that the information flow through an organization is usually separated from the material flow. When a lorry delivers an order to a customer, some information is included, but most is transmitted separately through the information system. The two are obviously connected, but managers have to control operations and check that the transactions reported actually happen – perhaps checking that a reported delivery actually arrives at a customer. One facet of these checks are periodic, manual counts of stock – called *stocktaking* – to find differences between recorded and actual stock levels. Any errors can be important when, say, a customer is promised immediate delivery of an item that is recorded as being

in stock, but there is actually no stock left. There are several reasons why actual stock may differ from recorded stock:

- records may not have been updated with recent transactions;
- stock is mislaid (particularly when an item is kept in several locations) and will turn up later;
- emergency or hurried withdrawals from stock may incur errors or not be recorded;
- deliveries or withdrawals are recorded more than once by mistake;
- cancelled orders or withdrawals are not recorded;
- poor quality and returned items (either from customers or to suppliers) are not recorded;
- stock become obsolete and is discarded;
- stock is stolen, damaged or deteriorates while in store.

Obviously, the match between actual and recorded stock should be as close as possible. There are several ways of encouraging this, and perhaps the most important is to restrict access to stores. Allowing only designated people into a store should make sure that there is no unauthorized use of the stock, and all transactions are properly recorded.

Many organizations feel that stocktaking is an expensive and inconvenient formality, and only do one check at the end of the financial year. Unfortunately, this means that discrepancies can last for a long time without being noticed. A more useful approach uses *cycle-counting,* where stock is checked at regular intervals. Typically, a small proportion of items is checked every week, with high usage items checked every month and less widely used items every three months. This counting can be expensive, so checks may be triggered by specific conditions such as a reported discrepancy in stock, an item with zero or low stock, an item with very high stock, a back-order placed for an item that is recorded as being in stock, or a period with no recorded demand.

Some discrepancies are inevitable in stocktaking, but how much is acceptable, and when does an organization have to do something about the errors? The answer depends entirely on the situation. If your stocks are barrels of whisky, there are surprisingly high losses through natural evaporation; if your stocks are gold bars you would expect no discrepancies at all. One common suggestion is that important items should have discrepancies of less than 0.2 per cent in stock levels, while other items should be within 1 per cent.

Summary

Inventory management depends on a lot of information from many sources. This is collected, organized, stored, analysed and presented by the inventory management information system. These systems come in many different forms, but they share

some common features. Most are based or a centralized store of information that is continuously updated by a transaction recording system.

Review questions

6.1 What is an inventory management information system?
6.2 What is the minimum amount of information needed by an inventory control system?
6.3 What are the main transactions in an inventory control system?
6.4 Where do the outputs from an inventory management information system go?
6.5 The more information a system can supply, the better the decisions made by management. Do you think this is true?

Information from accounting

Basic information

As we have already seen, some of the most important information for an inventory management system concerns the costs. This basic information is provided by accounting. At the very least, accountants supply data about the four cost components – unit, reorder, holding and shortage – but they also provide more detailed costs information and a lot of transactional data. This information appears in records of individual transactions, and is also summarized in the organization's accounts. Figure 6.3 shows a simple example of an entry in a set of trading accounts, where the profit is related to the change in stock value resulting from purchases and sales.

Trading account for month ending 31st January

	£	£		£
Opening stock		2,500	Sales	22,100
Purchases	13,000		Less returns	600
Less returns	500		Net sales	21,500
Net purchases	12,500			
Total stock		15,000		
Less closing stock		3,000		
Cost of products sold		12,000		
Gross Profit		9,500		
		21,500		21,500

Figure 6.3 Example of an entry in trading accounts

The 'cost of materials sold' is the total price paid, or cost of acquiring, the units that were later sold. Then the gross profit comes from:

$$\text{Gross profit} = \text{net sales revenue} - \text{cost of products sold}$$

In turn:

$$\text{Cost of products sold} = \text{opening stock} + \text{net purchases} - \text{closing stock}$$

As we saw in Chapter 2, there is an obvious problem in the way that we value remaining stock and, therefore, the cost of materials sold and profit. This problem remains when we go further with our account example in Figure 6.3, and find the cost of raw materials, work in progress and finished materials, as illustrated in Figure 6.4. We could continue to look at different parts of the accounts and see what information they give that is used by inventory managers, but the message is already clear. Inventory managers get a lot of information from accounting – and conversely, accounts get a lot of information from inventory management. We will limit ourselves to two illustrations of this. First, we will review the effect

Cost of products sold

Raw materials				
Opening stock of raw materials				$60,000
Purchases			$225,000	
Less	returns and allowances	$10,500		
	purchase discounts	$ 4,500	$ 15,000	
Net purchases			$210,000	
Add	inward transport		$ 15,000	
Cost of materials purchased				$225,000
Cost of materials available				$285,000
Less	closing stock of raw materials			$105,000
	Total cost of raw materials used			$180,000
Work in progress				
Raw materials used				$180,000
Staff				$390,000
Overheads				$125,000
Total cost of operations				$695,000
Add	opening stock of work in progress			$ 45,000
Total work in progress				$740,000
Less	closing stock of work in progress			$ 75,000
	Total cost of production			$665,000
Finished goods				
Opening stock of finished goods				$135,000
Add	cost of products			$665,000
Products available for sale				$800,000
Less	closing stock of finished goods			$180,000
	Total cost of products sold			$620,000

Figure 6.4 An example of stock value in product costing

mentioned in Chapter 2, where the profit depends on the value given to remaining stock. Second, we will do an ABC analysis, to see how the value of sales affects the effort put into inventory management.

Value of stock

An obvious problem with accounting information is that it depends on the conventions used. A change in accounting convention can markedly change the values used for inventory management, as we saw when valuing stock in Chapter 2.

Organizations need accurate values for their assets – including stock – as this directly affects their reported performance. Any errors can bring serious consequences. At times of high inflation, for example, the valuation of stocks is often too low and a company appears to have fewer assets than it really has. This may give an artificially high return on assets, and in extreme cases allows the company to be bought at a fraction of its true value.

A rigorous view says that stock has no real value until it is actually sold or used, and any value assigned before this is simply an accounting convenience. Nonetheless, organizations do need some reasonable view of their stock value, and the usual practice is to value stock at the lesser of unit cost or realizable value. Then, for each item:

$$\text{Value of stock} = \text{number of units in stock} \times \text{unit value}$$

If we add this for all items, we get the value of overall stock. Unfortunately, this is not so easy as both the stock level and the unit cost vary over time. We can avoid the problem of variable stock levels by counting the number of units held at a specific time, usually the end of the financial year. If there are enough items in the stock, this gives a reasonable overall average (providing the demands for items are more or less independent, and do not all follow the same strongly seasonal variation).

The other problem is variable unit cost. Inflation makes prices rise, and there are other factors, such as quantity discounts, different suppliers, variations in quality, different options or features, and different terms and trading conditions. As we saw in Chapter 2, there are four main options for dealing with this:

- *Actual cost* identifies each unit in stock with the price actually paid for it.
- *First-In-First-Out (FIFO)* This assumes that the first units arriving in stock are the first sold, so the value of remaining stock is set by the amount paid for the last units bought.
- *Last-In-First-Out (LIFO)* assumes that the latest units added to stock are used first, so the value of remaining stock is set by the amount paid for the earliest units bought.

- *Weighted average cost* finds the average unit cost over a typical period from:

$$\text{Average cost} = \frac{\text{Total cost of units}}{\text{Number of units bought}}$$

The method chosen should give the fairest match of costs to revenues. FIFO is the most widely used method, while LIFO is often expressly prohibited, especially when it would artificially reduce tax liabilities. Another consideration is that accounting conventions should be consistent, so that organizations should not look for a short-term gain by changing the basis of calculations.

Worked example

During her first job on a graduate trainee scheme, Madeleine Fraser was given the following record of transactions for an item and asked for her views on the profit. What answer could she give?

Date	Purchases Number	Unit cost	Sales number	Unit price
Opening stock	80	€20		
January			20	€30
February			40	€40
March	60	€30		
April			50	€50
May	80	€40		
June			20	€60

Solution

Madeleine has enough information to calculate the value of remaining stock, and hence the gross profit. At the end of June the closing stock is:

$$\begin{aligned}\text{closing stock} &= \text{opening stock} + \text{purchases} - \text{sales}\\ &= 80 + (60 + 80) - (20 + 40 + 50 + 20) = 90 \text{ units}\end{aligned}$$

The value of these units depends on the assumption used. She does not have enough information to use actual costs, but she can use the other three methods.

- *First-In-First-Out* assumes the 90 units remaining in stock are the last 90 bought, and this values stock at:

$$80 \times 40 + 10 \times 30 = £3{,}500 \text{ or } £38.89 \text{ a unit}$$

- *Last-In-First-Out* assumes the 90 units remaining in stock are the first units bought – after allowing for sales. Only 20 units of the opening stock can

remain, as the remaining 60 units were sold in January and February. Ten units of the stock bought in March can remain, and the remaining 60 units were bought in May. This values stock at: $20 \times 20 + 10 \times 30 + 60 \times 40 =$ £3,100 or £34.44 a unit

- *Weighted average cost* finds the average cost of all purchases as:

$$\frac{\text{total cost}}{\text{units purchased}} = \frac{80{\times}20 + 60{\times}30 + 80{\times}40}{80 + 60 + 80} = £30 \text{ a unit}$$

- This values stock at $90 \times 30 =$ £2,700

Now we know that:

gross profit = sales revenue − cost of units sold

sales revenue $= 20 \times 30 + 40 \times 40 + 50 \times 50 + 20 \times 60 =$ £5,900

cost of units sold = total cost of buying all units − present value of stock.

total cost of buying all units $= 80 \times 20 + 60 \times 30 + 80 \times 40 =$ £6,600

So:

Gross profit = 5,900 − [6,600 − present value of stock]

And with:

first-in-first-out:	gross profit = 5,900 − [6,600 − 3,500] = £2,800
last-in-first-out:	gross profit = 5,900 − [6,600 − 3,100] = £2,400
weighted average cost:	gross profit = 5,900 − [6,600 − 2,700] = £2,000

If inflation is high, the value of stock is actually increasing and adding to profit. In extreme cases, it might be more profitable not to sell the stock but to keep it and gain asset value (in the same way that house owners can become notionally rich when house prices rise quickly). On the other hand, any obsolete stock might be discarded by an organization as having no value, and then its assets and profits fall.

It is often expensive, difficult or impossible to count the actual number of units in stock. If you have stocks of, say, corn, iron ore, pencils or petrol, you may know approximately how much is in stock, but not the exact amount. We can, however, estimate the value of stock without actually counting it, by using the gross profit as a percentage of sales. If we assume this remains constant from one period to the next, we can work backwards to estimate the value of closing stocks.

Worked example

Last year PiHo Industries made a gross profit of 30 per cent of sales. This year they know that the opening stock of an item was $8,000 with purchases of $22,000. If sales of the item were $40,000, what is the value of closing stock?

Solution

If the gross profit is 30 per cent of sales, this comes to $12,000. The remaining 70 per cent, or $28,000, is the cost of products sold. We can use this in the following section of accounts:

Sales		$40,000	—from sales accounts
Cost of products sold:			
Opening stock	$8,000		= closing stock last period
Purchases	$22,000		—from purchase accounts
Cost of products available	$30,000		= purchases + opening stock
Estimated closing stock		?	—to be found
Total cost of products sold		$28,000	= 70% of sales
Gross profit		$12,000	= 30% of sales

The only value missing is the 'Estimated closing stock'. We know that the company had $30,000 of the item available, and it sold $28,000, so the difference must still be in stock. The estimated closing stock is (30,000 − 28,000 =) $2,000.

There are variations on this approach for estimating the value of closing stock, but they all rely on conventions and assumptions and are clearly less reliable than methods based on actual stock counts. They are most frequently used for checking values found from counts to see if they are reasonable, giving interim figures when it is too expensive or difficult to do a stock count, and when the actual stock has been lost through, say, fire.

ABC analysis of stocks

Another useful set of accounting information comes from an *ABC analysis*. Inventory control can take a lot of effort and so for some items, especially cheap ones, this effort is not worthwhile. Very few organizations, for example, include routine stationery or coffee in their formal stock systems. At the other end of the scale are very expensive items that need special care above the routine calculations. It would be useful to find the amount of effort worth putting into the control of any item. An ABC analysis gives some guidelines for this.

The origin of the ABC analysis – sometimes called Pareto analysis, or the 'rule of 80–20' – came in the nineteenth century when Vilfredo Pareto found that 20 per cent of the population owned 80 per cent of the wealth. In different forms, this is a widely used result, and in inventory control terms it means that 20 per cent of the inventory items need 80 per cent of the attention, while the remaining 80 per cent of items need 20 per cent of the attention. In particular, ABC analyses define the following:

A items are the few most expensive ones that need special care.

B items are ordinary ones that need standard care.

C items are the large number of cheap items that need little care.

Typically an organization might use an automated system to deal with all *B* items. *A* items are more important, and although the automated system might make some suggestions, managers make the final decisions after a thorough review of circumstances. *C* items are very cheap and are usually left out of the automatic system, to be dealt with by *ad hoc* procedures.

An ABC analysis starts by taking each item and multiplying the number of units used in a year by the unit cost. This gives the total annual use of items in terms of value. Usually, a few expensive items account for a lot of use, while many cheap ones account for little use. If we list the items in order of decreasing annual use by value, *A* items are at the top of the list and *C* items are at the bottom. We might typically find:

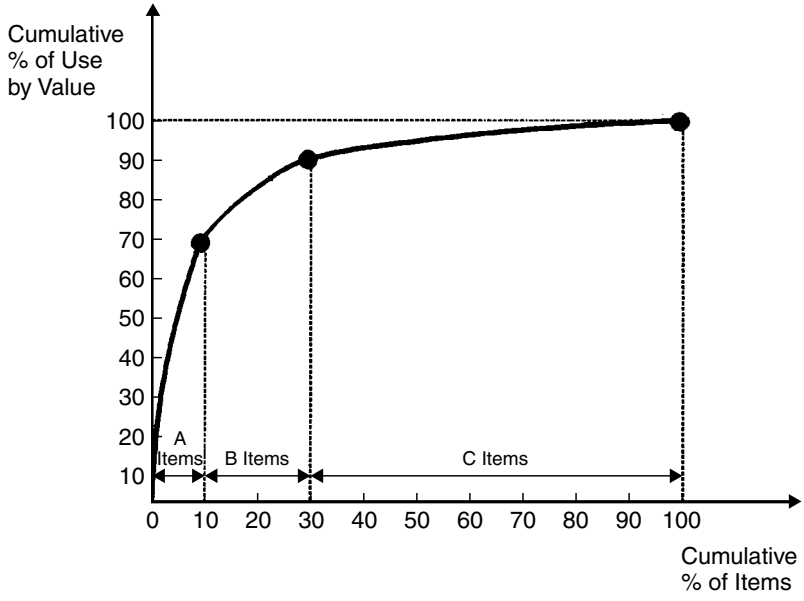


Figure 6.5 Typical ABC analysis of stock

Category	Percentage of items	Cumulative percentage of items	Percentage of use by value	Cumulative percentage of use by value
A	10	10	70	70
B	20	30	20	90
C	70	100	10	100

Plotting the cumulative percentage of annual use against the cumulative percentage of items gives the graph shown in Figure 6.5.

Worked example

A small store with 10 categories of item has the following costs and annual demands:

Item	X1	Y7	W4	X2	X3	Y9	W5	Z3	Z4	X4
Unit Cost (£)	3	2	3	8	2	10	1	5	20	4
Weekly demand ('00s)	2	25	1	30	10	10	5	2	1	3

Do an ABC analysis of these items. How might stocks of each category be controlled?

Solution

Although the demand figures are in hundreds of units a week, we only want comparisons and need not multiply each by 5,200 to give annual figures. So the weekly use of X1 in terms of value is $3 \times 2 = 6$, Y7 is 2×25, and so on. Repeating this calculation for the other items and sorting the results into order of decreasing annual use by value gives:

Item	X2	Y9	Y7	X3	Z4	X4	Z3	X1	W5	W4
Cumulative % of items	10	20	30	40	50	60	70	80	90	100
Weekly use	240	100	50	20	20	12	10	6	5	3
Cumulative use	240	340	390	410	430	442	452	458	463	466
Cumulative % of weekly use	52	73	84	88	92	95	97	98	99	100
Category	< A >	<-	-B-	->	<-	-	-	-	-	- - - - - - - - - - - - C - - - - - - - - - - - - >

The boundaries between categories are sometimes unclear, and we have to take fairly arbitrary decisions. In this case X2 is clearly an *A* item. The line between *B* and *C* items is probably best when Y9 and Y7 are *B* and all other items are *C*.

Item X2 accounts for 52 per cent of annual use, and these stocks should be carefully controlled, with managers making all relevant decisions. *B* items account for 36 per cent of use, and these can be controlled automatically. *C* items account for only 12 per cent of annual use and might be left to *ad hoc* procedures.

Having classified inventory items as *A, B* or *C*, you might think that *C* items are not important, and as well as putting no effort into their control you might not even bother stocking them. This would free resources to concentrate on the more important *A* and *B* items. In practice, this is usually a mistake. Because an item is cheap or has low demand, it does not mean that it is not important. A windscreen wiper is only a tiny part of the cost of a new car, but the car cannot be sold without one. A spare part for a production machine might not cost much and be rarely used, but if the machine breaks down, the shortage cost of not having the part in stock might be very high. The lesson is that organizations need stocks of all items – including those that are cheap or not often used – to continue their smooth operations. They typically hold *C* items in stock because they:

- are more important than their classification suggests – like spare parts;
- allow continued sales of an old item – which may no longer be made;
- are associated with sales of some *A* items – like windscreen wipers on a new car;
- give high profits in relation to their low costs;
- are new items – which have not yet established a history of use;
- are expected by customers – like ink cartridges for old printers.

If the three categories are not descriptive enough, managers might add extra ones. A fourth category is sometimes used, where *D* items are used so infrequently they are 'dead' and are being considered for withdrawal. And sometimes a special category is used for important spare parts or other items that have a combination of high importance and low use.

Summary

Accounting provides a lot of information for inventory management. This includes basic information about costs, and a range of transactional information about activities, suppliers and customers. We illustrated the use of accounting information for valuing stock, and showed how careful we have to be with the conventions used. ABC analyses categorize items according to use by value so that available effort can be sensibly shared out.

Review questions

6.6 Accountants provide a range of basic information that is essential for inventory management. Do you think this is true?
6.7 How can we find the absolute value of stock?
6.8 Is it LIFO or FIFO that generally undervalues stock?
6.9 What is the purpose of doing an ABC analyses of stock?
6.10 Which items are best dealt with by routine, automated control?

Information about supply and demand

Demand forecasting

The pattern of demand is perhaps the most important single factor for inventory management. After all, the purpose of stocks is to meet demands for items, so we must know what these demands are going to be. In general, there are two sources for demand data. For the first, an organization supplies products directly to customers and has historical sales figures that it can use to forecast future demand. This is such an important area that we will look at forecasting demand in the next chapter. For the second, an organization supplies materials to support internal operations and can use the planned operations to give a timetable for demands. Again, this is an important area that we look at in Chapters 8 and 9.

Supply and procurement

Inventory managers also need information about the pattern of supply. They have more control over this, through their procurement decisions. When an organization needs materials from suppliers, it initiates this supply with a purchase. Here the 'purchase' can include other types of transactions including rental, leasing, contracting, exchange, gifts, borrowing, producing, and so on. Because of these different options, some people prefer to talk about the 'acquisition of materials' or the more common term of *procurement*. Procurement is responsible for the flow of materials into stock.

- *Procurement* is responsible for acquiring all the materials needed by an organization.
- It consists of all the related activities needed to get materials, services and any other materials from suppliers into an organization.

The overall aim of procurement is to guarantee a reliable supply of materials to an organization. With this overriding aim, some more immediate goals include:

- working closely with user departments, developing relationships and understanding their needs;
- finding good suppliers, working closely with them and developing beneficial relationships;

- buying the right materials and making sure that they have acceptable quality, arrive at the time and place needed, and meet any other requirements;
- negotiating good prices and conditions;
- keeping stocks low, considering inventory policies, investment, standard and readily available materials, etc.;
- moving materials quickly through the supply chain, expediting deliveries when necessary;
- keeping abreast of conditions, including pending price rises, scarcities, new products, etc.

Procurement does not usually move materials itself, but it organizes the transfer. It gives the message that materials are needed, and arranges the change of ownership and location, but it is another function – such as transport – that actually delivers them. So procurement is largely concerned with information processing. It collects data from various sources, analyses it, and passes information onwards.

Procurement forms an essential link between organizations in the supply chain, and it gives a mechanism for co-ordinating the flow of materials between customers and suppliers; it passes messages backwards to describe what customers want, and it passes messages forwards to say what suppliers have available. Not only is procurement essential, but it is also responsible for a lot of expenditure – with 60 per cent of a typical manufacturer's costs spent on materials. A relatively small improvement in procurement can give substantial benefits. Suppose that a company buys raw materials for €60, spends €40 on operations and then sells the product for €110. It makes a profit of 110 – (60 + 40) = €10 a unit. Now suppose that procurement negotiates a 5 per cent discount on materials. Materials now cost $60 \times 0.95 = €57$, and with the same selling price the €3 saving goes straight to profit. The profit on each unit now jumps to €13, so a 5 per cent decrease in materials costs raises profit by 30 per cent.

Worked example

Last year Lavender Spartak Limited had total sales of $216 million. Their direct costs were $116 million for materials, $54 million for employees and $24 million for overheads. What is the effect of reducing the cost of materials by 1 per cent?

Solution

- The actual profit last year was $216 - (116 + 54 + 24) = \$22$ million.
- If the cost of materials drops by 1 per cent, it falls to $116 \times 0.99 = \$114.84$ million. Then the profit rises to $216 - (114.84 + 54 + 24) = \$23.16$ million. A 1 per cent decrease in materials costs increases profits by 5.3 per cent. Profit as a percentage of sales rises from 10.2 per cent to 10.7 per cent.

Procedure for procurement

Arguably, the most important part of procurement is finding the right supplier. There is no point in finding a good product or low prices, if the supplier cannot actually deliver. So procurement looks for two factors. First, they want materials that satisfy their requirements; second, they want a supplier who can guarantee to deliver the agreed materials. A useful approach for choosing the best supplier has the following steps.

- Look for alternative suppliers and build a long list of qualified suppliers who can deliver the materials.
- Compare organizations on this long list and eliminate those who are, for any reason, less desirable. Continue eliminating organizations until you have a shortlist of four or five of the most promising suppliers.
- Prepare an enquiry, or request for quotation, and send it to the shortlist.
- Collect the bids returned by the shortlist, do a preliminary evaluation, and eliminate those with major problems.
- Do a technical evaluation to see if the materials meet all specifications.
- Do a commercial evaluation to compare the costs and other conditions.
- Arrange a pre-award meeting to discuss bids with the remaining suppliers.
- At this point you should have enough information to make a final choice of supplier, so arrange a pre-commitment meeting to sort-out any last minute details.
- Award orders to the preferred supplier.

This procedure should give a supplier who is financially secure, with good long-term prospects, can supply the necessary materials, guarantees quality, is reliable, has short lead times, quotes reasonable prices and finance arrangements, is responsive to customers' needs, has necessary the skills and experience, has a good reputation, uses convenient procurement systems – and probably has many more enviable qualities. Now we can order materials, and again we need some prescribed procedure. This is different in every organization and varies with the type of item being purchased. We can, however, suggest a general approach. This has a series of common steps, which start with a user identifying a need for materials and end when the materials are delivered. A typical *procurement cycle* (outlined in Figure 6.6) has the following steps.

1. *A user department*:
   - identifies a need for purchased materials and prepares specifications;
   - checks budget and gets clearance to purchase;
   - prepares and sends a purchase request to procurement.

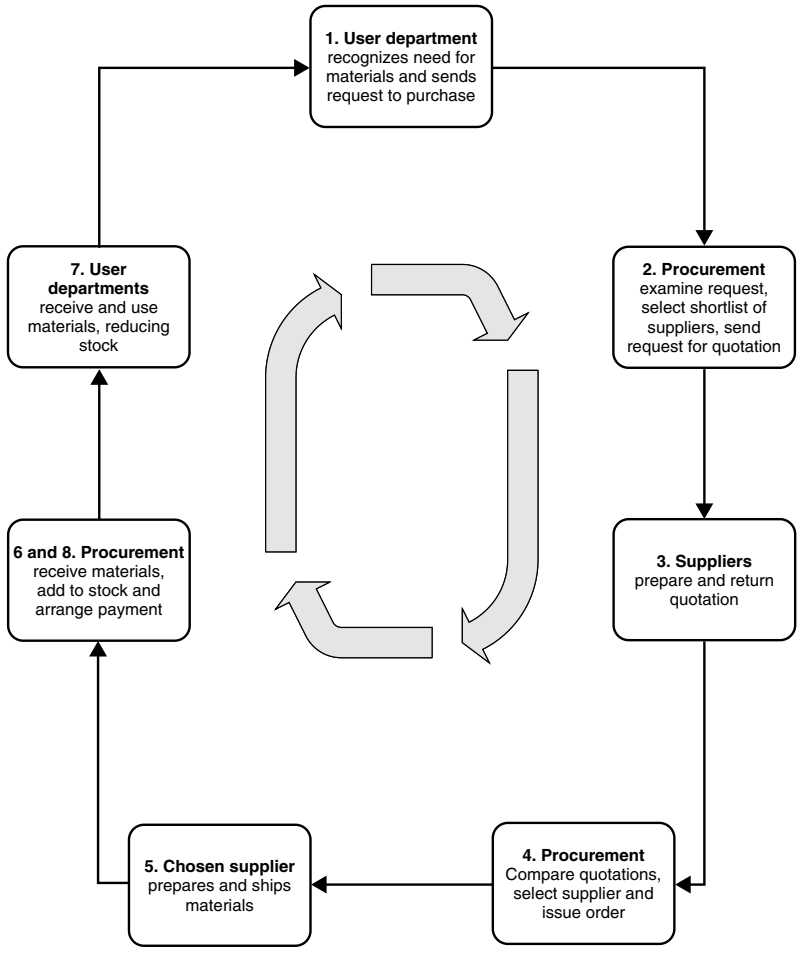

1. User department
recognizes need for materials and sends request to purchase
2. Procurement
examine request, select shortlist of suppliers, send request for quotation
3. Suppliers
prepare and return quotation
4. Procurement
Compare quotations, select supplier and issue order
5. Chosen supplier
prepares and ships materials
6 and 8. Procurement
receive materials, add to stock and arrange payment
7. User departments
receive and use materials, reducing stock

Figure 6.6 Structure of a purchasing cycle

2. *Then procurement*:
   - receives, verifies and checks the purchase request;
   - checks the material requested, looking at current stocks, alternative products, production options, etc. and then confirms the decision to purchase;
   - makes a shortlist of possible suppliers, from regular suppliers, lists of preferred suppliers, or those known to meet requirements;
   - sends a request for quotations to this shortlist.
3. *Then each supplier*:
   - examines the request for quotations;
   - checks the customer's status, credit, etc.;
   - sees how it can best satisfy the order and sends a quotation back to the organization, giving details of products, prices and conditions.

4. *Then procurement*:

   - examines the quotations and does commercial evaluations;
   - discusses technical aspects with the user department;
   - checks budget details and clearance to purchase;
   - chooses the best supplier, discuss, negotiate and finalize terms and conditions;
   - issues a purchase order for the materials, including terms and conditions.

5. *Then the chosen supplier*:

   - receives and processes the purchase order and organizes all operations needed to supply the materials;
   - ships materials together with a shipping advice and invoice.

6. *Then procurement*:

   - does any necessary follow-up and expediting;
   - acknowledges receipt, inspects and accepts the materials;
   - notifies the user department.

7. *Then the user department*:

   - receives and checks the materials, authorizes payment from budgets and updates inventory records;
   - uses the materials as needed.

8. *Then procurement*:

   - arranges payment of the supplier's invoice.

This procedure seems complicated, and involves many steps and documents. If you are buying something expensive – perhaps *A* items – this effort is certainly worthwhile and you may follow a much more complicated procedure to fix product specifications, select the supplier and negotiate terms. But if you are making small or routine purchases, have an existing relationship with a supplier, or there is only one qualified supplier, it is clearly not worth going through the whole procedure. Then you will look for more routine methods. A rule of thumb suggests that a routine order costs £80 to process, while Allen (2001) quotes figures of $115 to $150. For low value, *C*, items even routine procurement would cost more than the materials, and you should look for simpler, automatic or *ad hoc* procedures

Procurement is based on the analysis of large amounts of information. We can illustrate this by looking at step 5 of the cycle, when a supplier prepares and ships the material. To process the order the supplier needs information about the customer, product and stocks. If all goes well, and they can supply the materials from stock, they arrange with warehousing and dispatch to deliver the materials. If they have to replenish stocks, they have to plan operations and reorder materials from their own suppliers. This stage is illustrated in Figure 6.7, but this really only scratches the surface of the information flows.

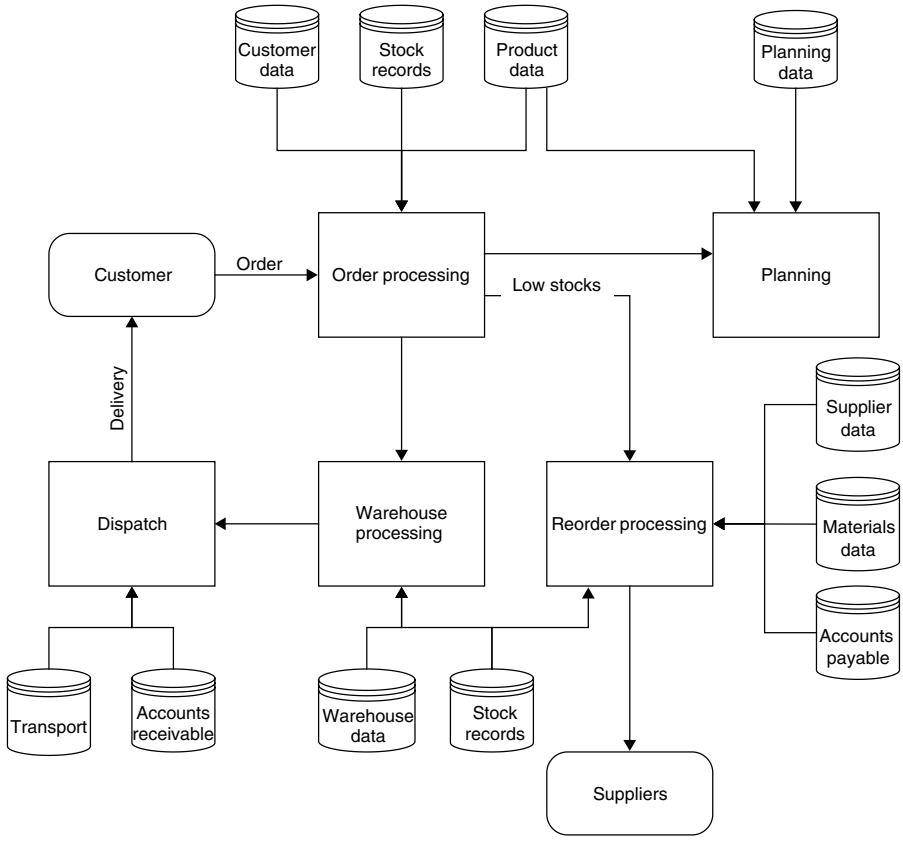


Figure 6.7 Some information flows for order processing

e-procurement

In recent years the status of procurement has risen considerably as it is recognized as an essential and expensive function. This has also been encouraged by changing patterns of procurement, with supply chains getting shorter, more customers use the Web, alliances are reducing the number of suppliers used by each organization, amounts purchased are increasing as companies focus on their core activities and outsource more, and customers are demanding more from products and conditions of purchase. Perhaps the most significant change in recent years has been the move to e-procurement. Surveys (see, for example, Cummings, 2002) suggest that over 60 per cent of UK companies were using e-procurement by 2002, and 80 per cent of European managers soon expect to use it extensively. Some of the advantages this brings are:

- allowing instant access to suppliers anywhere in the world;
- creating a transparent market where product details are readily available and easily compared;

- allowing comparison of standard terms of trade and rapid negotiation of details;
- automating procurement with standard procedures;
- greatly reducing the time needed for transactions;
- reducing costs, typically by 12–15 per cent;
- allowing outsourcing of some procurement activities to suppliers or third parties;
- integrating seamlessly with suppliers' information systems.

There are basically two types of e-procurement, which are described as B2B (where one business buys materials from another business) and B2C (when a final customer buys from a business). B2C has grown strongly, and between 1999 and 2002 the number of Internet shoppers in the UK rose from 2 million to 6 million (Rushe, 2001). Nonetheless, many sites hit financial troubles with the bursting of the 'dot.com bubble', and there have been a number of well-publicized bankruptcies. However, most e-procurement is actually B2B, with The Gartner Group (2001) giving the following estimates:

Year	Value of B2B procurement ($ billion)
1998	43
1999	219
2000	433
2001	919
2002	1900
2003	3600
2004	6000
2005	8500

The attractions of e-procurement are so great that most organizations will continue to move in this direction. The information flows can be much improved but people sometimes forget that e-procurement only improves the information flow and not necessarily the flow of actual materials. As Doerflinger *et al*. (1999) say, 'The real barrier to (B2B) entry is the back-end – fulfilment – not the Website itself'. Organizations clearly have to improve the other operations in the supply chain – including inventory management – if they want to take full advantage of e-business.

Summary

Procurement is the function responsible for the supply of items to an organization. It processes information about the supply of materials, from selecting suppliers

through to arranging payment for deliveries. This often follows a standard procedure in a purchasing cycle. Procurement has developed rapidly in recent years, particularly since the introduction of e-procurement.

Review questions

6.11 Why is procurement important to stock control?

6.12 Procurement moves materials from one organization to another. Do you think this is true?

6.13 There are so many advantages to e-procurement that soon all purchasing will be done electronically. When do you think that this will happen?

Warehousing

Purpose of warehouses

Most organizations keep their stocks in *warehouses*. In practice, these might be open fields where vegetables are stored, concrete pads for raw materials like coal and ores, sophisticated facilities that give the right conditions for frozen or delicate materials, databases that hold stocks of information, people who have a stock of skills, or almost any other form that you can think of. To put it simply, a warehouse is any place for storing materials. People use different terms for these, and when dealing with stocks of materials the most common are *distribution centres* and *logistics centres*.

In this broad sense, warehouses are an essential part of any supply chain, but Olsen (1996) comments that, 'We have seen the demise of warehousing predicted again and again, especially with the evolution of the philosophies of just-in-time, quick response, efficient consumer response, direct store delivery, and continuous flow distribution'. We will describe these methods in Chapter 10, but you should remember the reality that every organization keeps stocks – and as long as they keep stocks, they need warehouses to store them.

Warehousing and inventory management are so closely linked that they are sometimes considered to be the same function. In essence, though, inventory management is the management function associated with decisions about stock, while warehousing is the operational function that physically looks after it. Inventory managers decide what to put into stock, while warehousing physically receives the materials from suppliers and looks after them while they are in storage. Then we can summarize some of the main aims of a warehouse as providing necessary storage at key points in a supply chain, giving secure storage of the type needed by materials, keeping materials in good condition with little damage or loss, having low costs with high productivity and utilization of resources, and giving safe working conditions. To achieve these aims, the warehouse must perform a number of different activities, with the following list giving the most common:

- receiving materials from upstream suppliers;
- identifying the materials delivered, matching them to orders and finding their intended user;

- unloading materials from delivery vehicles;
- doing necessary checks on quantity, quality and condition;
- labelling materials (usually with bar codes or magnetic stripes) so they can be identified and monitored;
- sorting materials as needed;
- moving materials to a bulk storage area;
- holding them in stock until needed;
- when necessary, moving materials from bulk storage to a smaller picking store;
- picking materials from this store to meet orders;
- moving the materials to a marshalling area;
- assembling materials into orders;
- packing and packaging as necessary;
- loading delivery vehicles and dispatching the order;
- controlling all communications and related systems, such as inventory control and finance.

These activities might seem to reinforce the traditional role of warehouses as places for long-term storage of materials. In reality, though, organizations are moving materials quickly though supply chains, so warehouses are becoming staging points through which materials move as quickly as possible. They are increasingly seen as convenient locations to do a range of other jobs, such as sorting, packing and consolidating deliveries. They are, for example, places where small loads from different suppliers are combined to give full vehicle loads for delivery to customers. Or places where a manufacturer can combine parts made in different locations, such as the components for computer systems. This kind of consolidation can include final packaging to present a single product, or even do a limited amount of final manufacturing. This is the basis of postponement, where the final steps of production are left to the last possible moment.

Warehouses also do the opposite of consolidation when they break-bulk. Here a supplier sends all the demand for a particular area in a single delivery to a local warehouse. The warehouse breaks this delivery into the separate orders and passes them on to each customer.

In the extreme, warehouses do not put materials into storage at all, but are only transfer points. This is the approach recommended by Karabus and Croza (1995) who say that, 'product should never be warehoused or stored, but should continually be in movement, with the least possible number of handling steps'. This is the basis of *cross-docking*, which we mentioned in Chapter 1. The arrival of materials at a warehouse is co-ordinated with its departures to customers, so that they are transferred directly from the arrival area to the loading area, and immediately sent for delivery to downstream customers. As well as reducing stock

levels, this removes all the non-value adding activities of putting materials into storage, and later removing them.

Warehouse design

There are several important decisions for warehouses. As we saw in Chapter 2, these include the number and size of warehouses, and their locations. When these strategic decisions have been made, we can turn to more immediate decisions, such as the type of equipment and its layout. The list of activities above suggests the essential elements in a warehouse (illustrated in Figure 6.8) are:

- an arrival bay, or dock, where materials coming from suppliers are delivered, checked and sorted;
- a storage area, where materials are kept in stock;
- a departure bay, or dock, where customers' orders are assembled and sent out;
- a material handling system, for moving materials around;
- an information system, which records the location of all materials, arrivals from suppliers, departures to customers, and other relevant information.

Every time you go into a supermarket you see a sort of warehouse. Materials are delivered at the back of the supermarket, they are sorted and put onto shelves in the middle, then customers pick the items they want and take them away from the front. There are many variations on this basic outline. The most common one actually has two storage areas (as shown in Figure 6.9). When materials are delivered from suppliers they are put straight into a main bulk store. When required, packages in the bulk store are broken into individual units and moved to a smaller picking store. When a customer order arrives, the items needed are

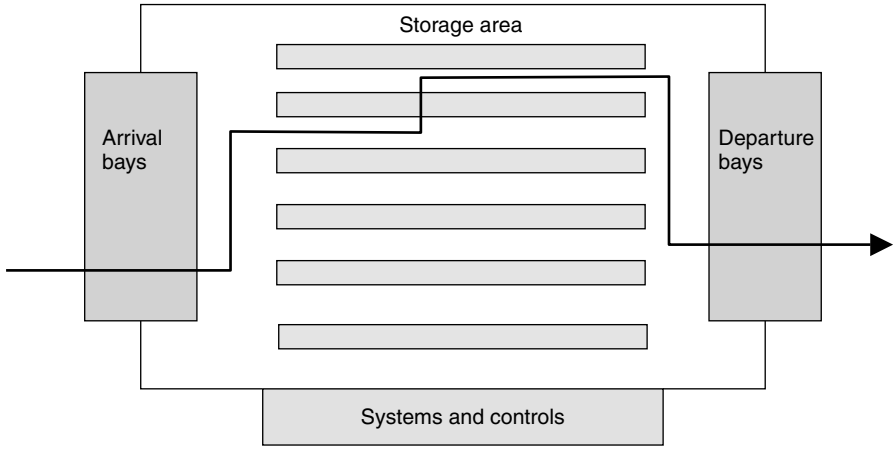


Figure 6.8 Basic layout of a warehouse

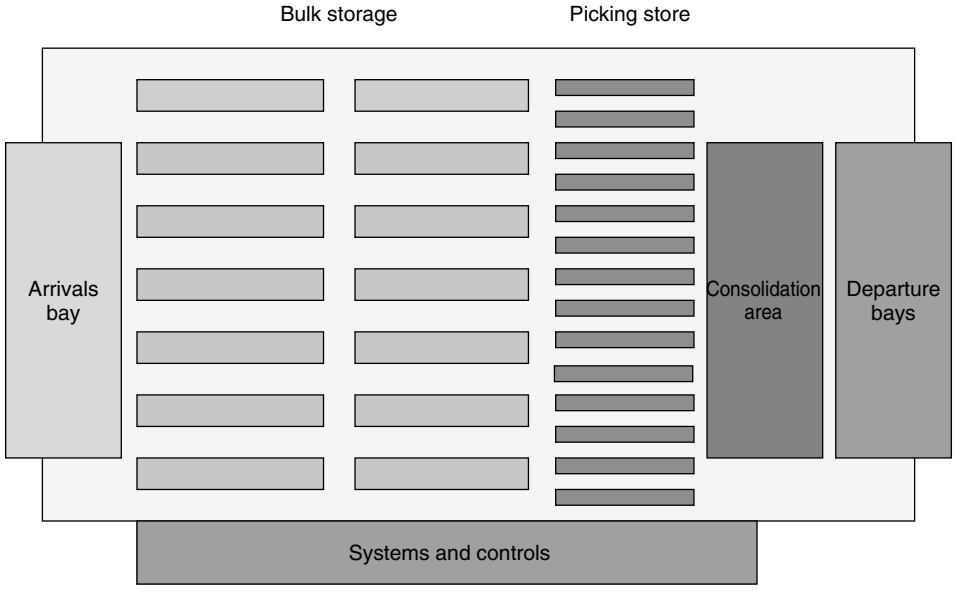


Figure 6.9 Schematic of a common warehouse layout

'picked' from the smaller, picking store and brought together in a consolidation area, before moving to the departure bays. When stocks in the picking store run low, they are replenished from the bulk store.

In practice, the overall layout is largely set by the shape of the existing building, the architect's views, the site available, height, or some physical constraint. Within these constraints, warehouse managers have to fit the details that best suit their operations. One way of approaching this detailed design is to do the following:

- estimate demand for materials over the next five years or so;
- translate this into forecast movements of materials into, through and out of the warehouse;
- compare available equipment for handling these movements and choose the best;
- find the space needed for storing and moving each item;
- design a general layout for the racking;
- see which materials should be close to each other (such as fast-moving materials nearer to transport bays, chilled materials in the same area, high value materials in safe areas, etc.) and those that should be far apart (such as foods far away from chemicals);
- develop outline plans for the facilities, compare alternatives and choose the best;
- add details to give final plans.

As you can see, the detailed layout depends on the type of storage and handling equipment. The most basic type of storage has an area of floor space, marked out in a grid to identify different locations. Bulky or heavy items are put into a location, probably by a forklift truck. The next level of storage uses shelves or racking built in aisles, with materials typically on pallets. To reduce the ground area, these aisles can be quite high. But the shelves have to be shallow, so that all materials are within reach, and the aisles tend to be long. The typical design has racking that is long, high and narrow. Flow racks can increase the density of storage by making the shelves much deeper. These are sloping shelves that are filled from the back, and as you remove a unit from the front, all the remaining units roll forward. Other options for storage include pigeonholes for small units, horizontal carousels (bins on an oval track that rotate to bring materials to a picker), vertical carousels (shelves that rotate up and down to bring materials within reach), hanging racks for garments, silos and tanks for fluids, and a huge assortment of other arrangements.

Related to the storage of materials is the handling equipment for moving them, and there are three main alternatives.

1. *Manual warehouses*. This is probably the easiest arrangement to imagine, with items stored on shelves or in bins. People go around, pick items from the shelves and move them around in some sort of container, like a supermarket trolley. Other equipment for moving materials includes hand trucks for moving pallets, and carousels to bring materials to pickers, but essentially people control all movements. Manual warehouses only work if the items are small and light enough to lift, with shelves that are low enough to reach and close together to reduce walking.
2. *Mechanized warehouses* replace some of the muscle power of manual warehouses by machines. Typical examples of mechanized equipment are:
   - reach trucks, which are usually electrically powered and move pallets and other loads up to storage locations. A driver stays on the ground to control the truck, which can raise loads vertically to a considerable height.
   - order-picking machines are a variation on reach trucks, where the driver is lifted with the materials to pick, or deliver, at high locations.
   - forklift trucks, come in many different versions and are the standard means of moving pallets and equivalent loads for short distances. They are very manoeuvrable and flexible, but they need space to work and are fairly expensive to use.
   - cranes, which describe a family of vehicles used to lift materials;
   - towlines, which are continuous cables that can move trailers around a fixed path, rather like ski lifts;
   - conveyors, which move large quantities of materials along fixed paths;
   - tractors or trains, which are power units that pull loads on trailer units, rather like small articulated lorries;
   - carousels, which are a continuous series of bins going round a fixed track. At some point on the journey items are put into a bin, and the bins are emptied when they pass another chute or collection point.

The key point about mechanized systems is that they are still under the control of an operator. Someone actually drives a forklift or controls the movement along a towline. The next alternative is to pass the control of movements to a computer.

3. *Automated warehouses*. Traditional warehouses, even mechanized ones, tend to have high operating costs. These operating costs can be reduced, as well as improving aspects of service, by using automation. Unfortunately, this needs a heavy investment in equipment, and is only worthwhile for very big stores. Typical elements of an automated warehouse are:
   - storage areas that can be accessed by automatic equipment; these often use narrow aisles up to, say, 40 m tall to get a high density of materials and minimize the distances moved;
   - equipment to move materials around the warehouse; these are usually automated guided vehicles (AGVs) which are controlled by wires in the floor, but might include conveyors, tractors, or a range of other moving equipment;
   - equipment automatically to pick materials and put them into storage, including high speed stacker cranes that can reach any point in the narrow aisles very quickly;
   - equipment to transfer materials between the different types of equipment; these automatic loaders and unloaders might include industrial robots;
   - a warehouse management system to record material locations, and control all movements.

Packaging

Many materials are moved on pallets (standard wooden trays about four feet square) and in containers (the 20- or 40-foot metal boxes). Collecting together materials into these standard packages is called *unitization* to form *unit loads*. If an organization uses standard loads, it can set up all its handling equipment to move these efficiently, rather than have to deal with all sorts of different shapes and sizes.

Even if they are put into unit loads, many items – particularly delicate items such as china and electronics – need special protective packaging. Sometimes the packaging can protect materials from harsh environments, such as rain or sun; sometimes it is needed to separate materials that would contaminate each other, such as sugar and petrol; sometime it keeps the contents clean, such as foodstuffs and medicines. In general, packaging serves four basic functions as it:

- identifies the product and gives basic information;
- protects items while they are being moved through the supply chain;
- makes handling easier;
- assists in marketing, promoting the product, advertising and giving information to customers.

The balance between these depends very much on the product. Bars of chocolate, for example, might put more emphasis on marketing, while boxes of ice cream

might be more concerned with protection. There are also two types of packaging to consider. First, the interior, or consumer packaging, is designed for the customer and includes the marketing and promotional materials. Second, the exterior, or industrial packaging, is designed to identify the contents, give information to organizations in the supply chain, protect the contents, and make handling easier. You probably think of packaging in terms of cardboard boxes, but it can be made of many different materials. Consumer packaging tends towards bright colours, plastic and cellophane, while industrial packaging tends towards dull cardboard. Some perfumes and alcoholic drinks come in very elaborate containers that cost more to produce than the contents. This is an area where reuse and recycling are becoming more important, and legislation is reducing the amount of packaging that is simply discarded.

Summary

Warehousing is the function that is responsible for the physical storage of materials. It consists of a range of activities concerned with the storage, movement, identification, sorting, packaging, etc. of materials. Strategic decisions set the size and location of warehouses; more immediate decisions concern the layout, type of handling equipment and packaging needed.

Review questions

6.14 What is the difference between warehousing and inventory management?

6.15 What are the main decisions needed for warehouses?

6.16 What are the three levels of technology found in warehouses?

6.17 What is the purpose of packaging?

Chapter review

- This chapter has described some of the information needed for inventory management.
- Inventory management depends on a lot of information from many sources. This is collected, organized, stored, analysed and presented by the inventory management information system. These systems come in many different forms, but they share some common features. Most are based around a centralized store of information that is continuously updated by a transaction recording system.
- Accounting provides a lot of information for inventory management. This includes basic information about costs, and a range of transactional information about activities, suppliers and customers. We illustrated the use of accounting information for valuing stock, and showed the importance of being careful with conventions.
- ABC analyses categorize items according to use by value so that available effort can be sensibly shared out. *A* items are most expensive and should be given special care; *C* items are cheapest and need little control.

- Procurement is the function responsible for the supply of items to an organization. It processes information about the supply of materials, from selecting suppliers through to arranging payment for deliveries. This often follows a standard procedure in a purchasing cycle. Procurement has developed rapidly in recent years, particularly since the introduction of e-procurement.
- Warehousing is responsible for the physical storage of materials. It consists of a range of activities concerned with the storage, movement, identification, sorting, packaging, etc. of materials.
- Strategic decisions set the size and location of warehouses; more immediate decisions concern the layout, type of handling equipment and packaging needed.

Project

If you look at the detailed accounts of a company, you can get a lot of information about their stocks. The aim of this project, then, is to examine the accounts published by a major company. If you do not have a hard copy of the accounts, you can find them on the company's website, or websites sponsored by financial commentators. See how stocks are represented in the accounts and describe the information about that you can extract. What other information would you like to make sensible decisions about inventory policy?

Problems

6.1 Julie Atherstone paid €100 a unit for items which she sold for €200. The last delivery of 100 units arrived when there were 100 units left in stock, and the supplier told Julie that the unit cost had risen to €125. She immediately raised her price to €250, and at the end of the financial year she had sold 140 units. Can you say anything about Julie's financial performance?

6.2 Accounts provided by Netherton Hall include the following transactions for an item.

Date	Purchases number	Unit cost	Sales number	Unit price
Jan 1 – Opening Stock	120	£21		
Jan			60	£32
Feb			50	£37
Mar	180	£26	20	£37
Apr			110	£39
May	90	£28		
Jun			80	£39
Jul	20	£28		

What is the value of stock held at the end of July and the gross profit from sales?

6.3 Karl Ehrenberg was asked to find the value of ore stocked by his company. He knows that last year's gross profit was 25 per cent of sales. Opening stock of the ore was €70,000 with purchases of €300,000 and sales of €400,000. How could Karl estimate the current stock?

6.4 A store wants to improve the control of its stocks, and is looking at the possibility of using a Pareto analysis. Records from eight types of item show the current sales and costs as follows.

Item	3427	1822	5362	2777	1413	9719	5520	0188
Number of sales	25	150	30	80	10	40	1,000	100
Cost ($)	1,400	14	680	20	1,020	150	20	30

Do a Pareto analysis on these items.

6.5 Belgrave Furnishing Retail Ltd has 20 categories of furniture with the following costs and annual demands:

Item	0	1	2	3	4	5	6	7	8	9
Unit cost (£)	40	20	40	100	20	100	10	40	200	10
Weekly demand	5	100	40	30	30	10	20	10	5	100
Item	10	11	12	13	14	15	16	17	18	19
Unit cost (£)	50	30	65	150	40	180	25	80	100	5
Weekly demand	10	120	50	30	40	15	25	20	5	120

Do an ABC analysis for Belgrave. If resources are limited, which items should they give most attention to?

Discussion questions

6.1 Inventory managers need a lot of information but they can get swamped by having too much and get 'information overload'. What information is really needed for inventory management, and how can you tell if there is too much?

6.2 An oil company buys crude oil on the open market, transports and refines it, and sells petrol at a chain of filling stations. The price of crude oil fluctuates quite widely. What policy is the company likely to use for valuing its stock? Can you imagine managers manipulating stock holding costs to make themselves appear in the best light?

6.3 The most important information for inventory management comes from accounts, and this depends on the conventions that are used. Changing

the conventions can give major changes in perceived costs and consequent inventory policy. Is this true? And if it is, how can we overcome these effects?

6.4 One of the most important trends in recent years has been towards e-commerce. How does this affect inventory management?

6.5 In recent years purchasing has moved from a largely clerical job to a higher status profession. Can you find any evidence for this change, or identify any likely causes?

6.6 What is the difference between inventory management and warehousing? If stock is managed properly, do we still need warehouses?

References and further reading

Allen, S. (2001) Leveraging procurement: the quiet e-business, *Logistics and Transport Focus*, 3(4), 29–30.

Bodnar, G. and Hopwood, W.S. (2000) *Accounting Information Systems* (8th edition), Englewood Cliffs, NJ: Prentice Hall.

Chaffey, D., Bocij, P., Greasley, A. and Hickie, S. (2002) *Business Information Systems*, Englewood Cliffs, NJ: Prentice Hall.

Cummings, N. (2002) UK leading the world in e-procurement, *OR Newsletter*, March, 9.

Curtis, G. and Cobham, D. (2001) *Business Information Systems* (4th edition), Harlow: FT-Prentice Hall.

Dauderis, H. (1990) *Financial Accounting*, Toronto: Holt, Rinehart and Winston.

Doerflinger, T.M., Gerharty, M. and Kerschner, E.M. (1999) The information revolution wars, *Paine-Webber Newsletter*, New York.

Karabus, A. and Croza, M. (1995) The keys to the kingdom, *Materials Management and Distribution*, May, 21–22.

Olsen, D.R. (1996) Warehousing trends for the next generation, *Logistics Focus*, 4(2), 6–8.

Romney, M.B. and Steinbart, P.J. (2002) *Accounting Information Systems*, Englewood Cliffs, NJ: Prentice Hall.

Rushe, D. (2001) WWW.basketcase, *The Sunday Times*, 2 September, 5.

Waters, D. (ed.) (2003) *Global Logistics and Distribution Planning*, London: Kogan Page.

7

Forecasting Demand

Aims of the chapter

In the last chapter we looked at the information needed for inventory management. We developed some general ideas, and looked specifically at the information supplied by accounting, procurement and warehousing. We mentioned two other important sources of information – forecasts of demand and plans for future operations. In this chapter we discuss methods of forecasting demand. In the next chapter we look at the planning of operations.

After reading this chapter you should be able to the following:

- discuss the role of forecasting in inventory management;
- review different approaches to forecasting;
- use a variety of judgemental forecasting methods;
- define 'time series' and appreciate their importance for inventory control;
- calculate forecast errors;
- describe the characteristics of causal forecasting and use linear regression;
- describe the characteristics of projective forecasting and use forecasts based on simple averages, moving averages and exponential smoothing;
- forecasts demand with seasonality and trend;
- consider the planning needed for forecasts.

This chapter will emphasize:

- *forecasts*, which predict future circumstances;
- *judgemental forecasts*, which are based on opinion;
- *causal forecasts*, which use relationships between variables to forecast;
- *projective forecasts*, which project past patterns into the future.

Methods of forecasting

Use of forecasts

Chapters 3 to 5 described some quantitative models for inventory control. The most important piece of information for these models is often the shape of future demand, and with independent demands we have to find this from forecasts. The problem of forecasting values for the future is not, of course, unique to inventory control. All business plans become effective at some point in the future, so they must be based not on current circumstances, but on the prevailing circumstances when the plans become effective. We cannot know with certainty what will happen in the future, but the best we can do is forecast expected conditions. Then if our forecasts are accurate, our plans are based on the right information – but if our forecasts are wrong, our plans are based on faulty assumptions and we inevitably get poor results.

Figure 7.1 shows how forecasts can fit into inventory control. Managers review a range of information that they have about objectives, constraints, alternatives, past performance, the environment, etc., and then look at the forecasts for future conditions and make their decisions. Later, details of actual performance are reviewed, errors noted, forecasts updated and so the cycle continues. As you can see, forecasting is never finished but has continuous updates and revisions. You

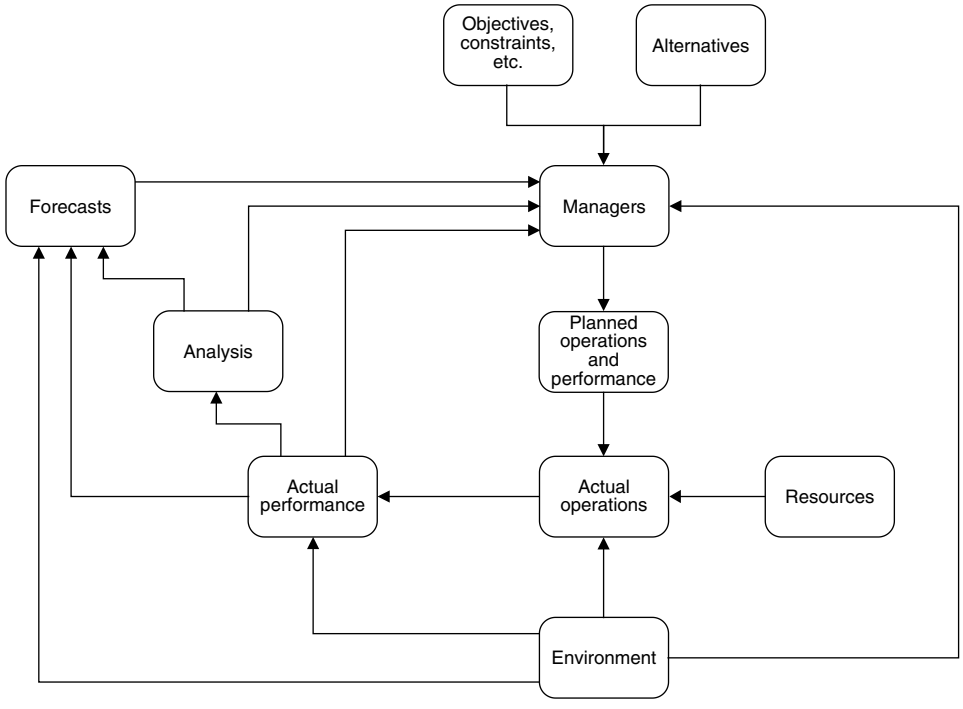


Figure 7.1 Forecasting supporting management decisions

can also see that forecasts are not prepared by a group of isolated specialists, but are prepared and used throughout the whole organization.

Here we emphasize the forecasting of demand, but remember that almost everything has to forecast to some extent. When, for example, we use cost figures in an inventory control model, these are really forecasts of the costs that will apply in the future when it is time to place an order. Often we take a very informal approach to these forecasts and assume that costs will remain constant for the immediate future. Another important area to forecast is the lead time. Lead times are generally getting shorter, so we should remember this when looking at the lead time demand. We should also remember that if the lead times become short enough we may not have to forecast customer demand at all (Thacker, 2001a). This happens when an organization's customers are prepared to wait for longer than the lead time given by suppliers to the organization. Then the organization need hold no stock, but only pass on the demand when it gets an order from a customer.

Forecasts give important inputs to inventory management but, in turn, they need information from other sources. This information includes the best type of forecasting model, values for parameters, historical data, subjective inputs, and so on. Figure 7.2 outlines the flow of information for a forecast. This has an initial forecast produced automatically, and then managers review the result and make adjustments (based on their experience, knowledge and other information) to give the final forecast. In practice, even a small inventory contains thousands of items, so managers do not have much time for this review. The forecasting routines

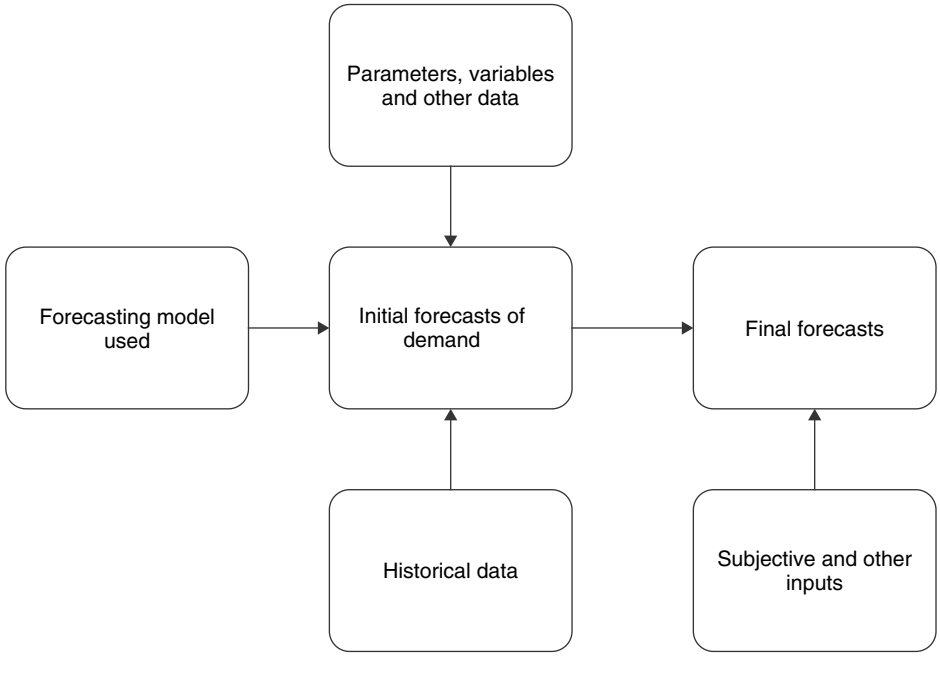


Figure 7.2 Information needed to produce forecasts

should, therefore, automatically produce the type of results that managers want, even if these are not the theoretical optimum. These forecasts should be:

- *accurate* – with small errors
- *unbiased* – so they do not always under- or over-estimate demand
- *responsive* to changes in demand
- *not affected* by the odd unusual figure
- *in time* for its purpose
- *cost-effective*
- *easy to understand*.

They should also recognize the specific areas where management input is needed and produce an 'exception report' for items that are particularly important or expensive, have large or erratic forecasting errors, have demand that suddenly changes, have a major change of some other type, have no recent demand, or have recently been introduced to stock.

Methods of forecasting

There are so many different ways of forecasting, so many different things to forecast and so many different circumstances, that no single method of forecasting is always the best. We have to choose a method that suits our needs. This choice depends on many factors, including the following:

- time covered in the future
- availability of historical data
- relevance of historical data to the future
- type of product
- variability of demand
- accuracy needed and cost of errors
- benefits expected from the forecasts
- amount of money and time available for the forecast.

You might think that the most complex and expensive methods give best results but this is not necessarily true and simple methods can give very good results, while complex methods can give very poor ones. We should, therefore, look carefully at the choice of method and not jump at the most expensive. But what are the methods available? We can classify them in several ways, starting with the time covered in the future.

1. *Long-term forecasts* look ahead several years – the time needed to build a new factory or organize new facilities. They usually look at overall demand, so a brewery might forecast total demand of 10,000 barrels a day in three years' time. This gives enough information to plan budgets and major facilities over the next few years.
2. *Medium-term* forecasts look ahead between three months and a year – the time needed to replace an old product by a new one or organize resources. The brewery above, for example, might forecast 1,000 barrels a day of bitter, over the next year, and use this to plan production and associated facilities.
3. *Short-term* forecasts cover the next few weeks – describing the continuing demand for a product or scheduling operations. Again, the brewery might forecast 10 barrels of premium bitter needed by the Red Star Pub, and use this to schedule deliveries and stock movements.

The time horizon affects the choice of forecasting method, because of the availability and relevance of historic data, the time available to do the forecasting, the cost involved and the effort considered worthwhile. In essence, long-term forecasts are concerned with strategic decisions, medium-term forecasts with tactical decisions, and short-term forecasts with operational decisions. Inventory control is primarily concerned with short-term forecasting, so there is usually enough relevant data to make a reasonable forecast.

Another classification of methods shows the difference between qualitative and quantitative forecasts (shown in Figure 7.3). If an organization is already making a product, it has records of past demand and knows the factors that affect this. Then it can use a quantitative method to forecast future demand. There are two alternative approaches:

1. *Projective methods* look at the pattern of past demand and extend this into the future. If demand in the last 4 weeks has been 100, 110, 120 and 130 units, we can project this pattern and suggest that next week's demand will be around 140 units.
2. *Causal methods* look at the factors that affect demand and use these to forecast. The number of units of a product sold depends on the price charged, so we can use the planned price to forecast sales.

Both of these approaches rely on accurate, numerical data. Suppose, though, that an organization is introducing an entirely new product. There are obviously no past demand figures to project into the future, and the organization does not yet know what factors affect demand. As there is no data for a quantitative method, the only alternative is a qualitative one. These are generally called *judgemental*, and they rely on subjective views and opinions.

Summary

All decisions about stocks become effective at some point in the future and depend on forecasts of prevailing conditions. There are many methods of forecasting, but

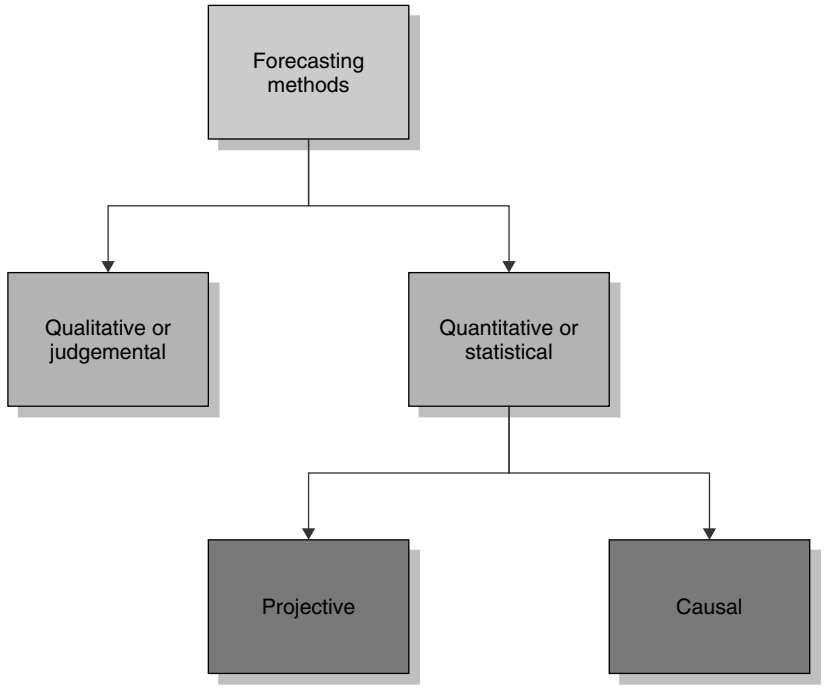


Figure 7.3 Approaches to forecasting

none is always the best. We can classify the methods in several ways, with useful ones describing the time they look ahead (in the short, medium or long term), and the overall approach (judgement, causal or projective).

Review questions

7.1 Why is forecasting important for inventory management?
7.2 List three different approaches to forecasting.
7.3 What factors should be considered when choosing a forecasting method?
7.4 Forecasting is a specialized function that uses sophisticated methods to project historical trends into the future. Do you think this is true?

Judgemental forecasts

Judgemental forecasting methods are subjective assessments, usually based on the opinions of experts. For inventory management these experts might include suppliers, purchasing departments, store keepers, salesmen, customers, organizations supplying similar or related items, trade reviews, government publications, and so on. The methods are very flexible and can be used in a wide range of circumstances, but are not as reliable as quantitative forecasts. If, however,

an organization is about to stock an entirely new item, it has no historical data for a quantitative method and it must use a judgemental forecast. Sometimes, even when there is data, it is unreliable, out of date, or irrelevant to the future. Then the organization has no alternative but to use a judgemental method. The five most widely used methods of judgemental forecasting are personal insight, panel consensus, market surveys, historical analogy and the Delphi method.

1. *Personal insight*. This uses a single expert who is familiar with the situation to produce a forecast based on his or her own judgement. This is the most widely used forecasting method – and is the one that managers should try to avoid. It relies entirely on one person's judgement – as well as their opinions, bias, prejudices, ignorance and mood. Comparisons of forecasting methods clearly show that someone who is familiar with a situation and uses experience and subjective opinions to forecast will consistently produce *worse* forecasts than someone who knows nothing about the situation but uses a more formal method.

Personal insight can give good forecasts, but often gives very bad ones and there are countless examples of experts being totally wrong. So the major weakness of the method is its unreliability. This may not matter for minor decisions, but when the consequences of errors are important, you should use a more reliable method.

2. *Panel consensus*. One person can make a mistake, but we should get better results by collecting together several experts and allowing them to talk freely to each other until they reach a consensus. When there is no secrecy and the panel are encouraged to talk openly, a genuine consensus can emerge. On the other hand, it can be difficult to get the panel to talk openly, or to combine their different views into a consensus.

Although it is more reliable than personal insight, panel consensus still has the major weakness that everyone – even experts – can make mistakes and the consensus might still be wrong. There are also problems of group working, where 'those who shout loudest get their way', everyone tries to please the boss, some people do not speak well in groups, and so on. Overall, panel consensus is an improvement on personal insight, but you should view results from both methods with caution.

3. *Market surveys*. Sometimes it is better to ask the people most closely concerned rather than get the opinion of experts. When launching a new product, for example, it might be better to get the opinions of potential customers. This is the basis of market surveys, which collect data from a representative sample of customers, analyse their views, and then draw inferences about the population at large. This tends to be expensive and time-consuming, but it can get good results. You can, however, still find market surveys that have given very poor results, largely because they put too much reliance on:

- a sample of customers that accurately represents the population;
- carefully worded, useful, unbiased questions;

- fair and honest answers;
- reliable analyses of the answers;
- valid conclusions drawn from the analyses.

4. *Historical analogy*. Most products have a finite lifespan, and during this the demand follows a common pattern with periods of introduction, growth, maturity, decline and withdrawal (as illustrated in Figure 7.4). Historical analogy uses the demand of a similar item that was introduced in the past to judge the demand for a new item. A publisher, for example, forecasts sales of a new book by assuming they follow the same pattern as a similar book that it published recently. The main problems, of course, are finding a recently introduced product that is similar enough, and fitting the characteristic life-cycle curve to actual demand.

5. *Delphi method*. This is the most formal of the judgemental methods and has a well-defined procedure. A number of experts are posted a questionnaire to ask their opinions. This questionnaire avoids the problems of panel consensus, and each reply is anonymous to avoid the influences of status, etc. The replies from these questionnaires are analysed and summaries are passed back to the experts. Now each expert is asked to reconsider their original reply in the light of the summarized replies from others. They may be convinced by some of the arguments, and adjust their answers for a second round of opinions. This process of modifying responses in the light of replies made by the rest of the group is repeated several times – usually between three and six. By this time, the range of opinions should be narrow enough to help with decisions.

The main problems with the Delphi method are designing appropriate questionnaires, finding a suitable mix of experts, the time involved, and keeping

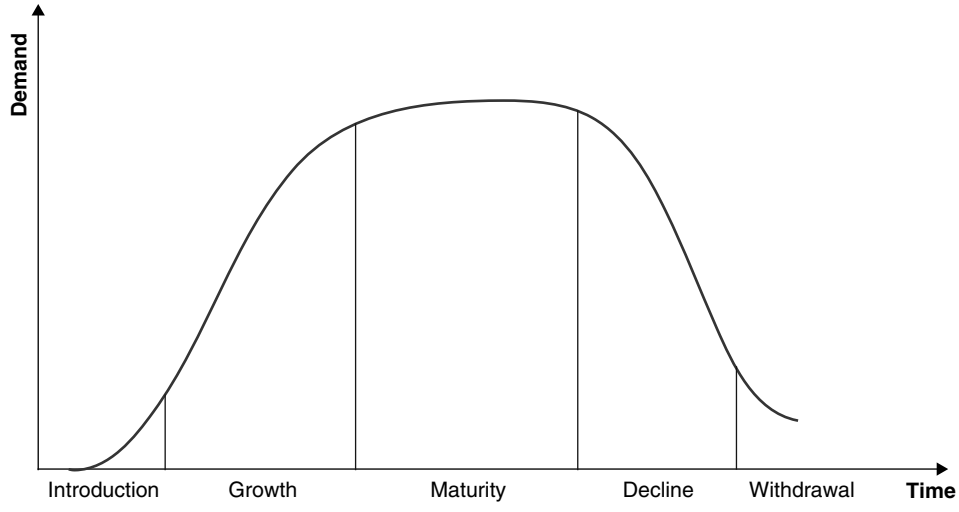


Figure 7.4 Demand at different stages of a life cycle

the same group involved over this time. The experts are also likely to give answers that depend on their own responsibilities and aims rather than objective analyses. A lot of anecdotal evidence suggests that these surveys giving disappointing results.

Each of the judgemental methods works best in different circumstances. If you want a quick reply, personal insight is the fastest and cheapest method. For many inventories, suppliers are the best source of information as they have knowledge and experience with similar organizations. If you want more reliable forecasts, it may be worth organizing a market survey or Delphi method.

Summary

Judgemental forecasts rely on subjective views and opinions of experts. They are the only possible methods when there is no reliable historical data. The most common methods are personal insight, panel consensus, market surveys, historical analogy and Delphi method.

Review questions

7.5 When are judgemental forecasts most likely to be used?
7.6 List five types of judgemental forecast.
7.7 What are the main problems with judgemental forecasts?

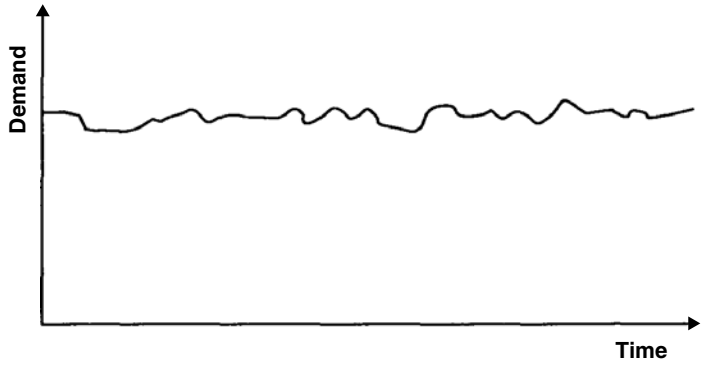
Time series

Quantitative forecasts are often based on *time series*, which are series of observations taken at regular intervals of time. The weekly demand for a product, monthly unemployment figures, daily rainfall and annual population statistics are examples of time series.

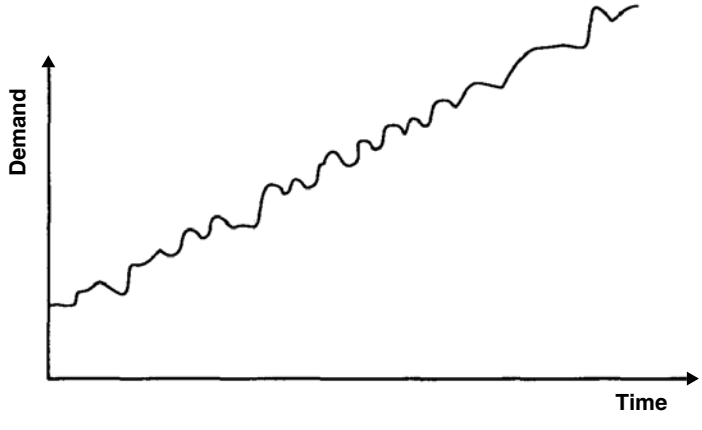
If you have a time series, the best way to start analysing it is to draw a graph. You can see the underlying patterns in this, with the three most common (shown in Figure 7.5):

- *constant series*, where demand continues at roughly the same level over time (such as demand for bread or annual rainfall);
- *trend*, where demand either rises or falls steadily (such as demand for 3G phones or the price of petrol);
- *seasonality*, where demand has a cyclical component (such as demand for ice cream or electricity).

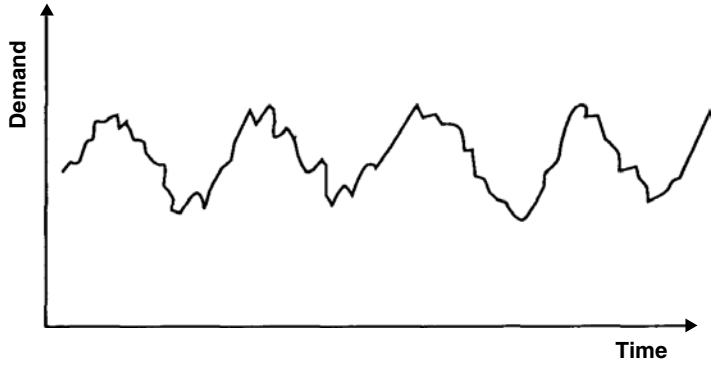
Forecasting would be easy if demand followed such simple patterns. Unfortunately, there are always differences between actual demand and the underlying



(a) Constant series



(b) Trend



(c) Seasonality

Figure 7.5 Common patterns in time series

pattern. These differences form a random *noise* that is superimposed on the underlying pattern. A constant series, for example, does not always take exactly the same value, and:

101 105 93 96 108 103 98 101 91 95 109

is a constant series of 100 with superimposed noise.

$$\text{Actual demand} = \text{underlying pattern} + \text{random noise}$$

The noise is a completely random effect that is caused by many factors, such as varying customer demand, hours worked, speed of working, weather, rejections at inspections, time of year, wider economic influences, errors in available data, delays in updating information, poor communications, and so on. It is the noise that makes forecasting difficult. If the noise is relatively small, actual demand is close to the underlying pattern and we can get good forecasts: if there is a lot of noise, it hides the underlying pattern and forecasting is more difficult.

Because of the noise, our forecasts are almost always wrong. In other words, there is a difference between the forecast and the actual values. If we forecast next week's demand for a product as 20 units, we would not be surprised if actual demand turns out to be 22 units. Then our forecast has an *error* of:

$$\text{error} = \text{actual demand} - \text{forecast} = 22 - 20 = 2 \text{ units}$$

With a good forecast this error should be relatively small. If we continue making a forecast every period, we can see how it performs over time. We could, for example, find the error every period and use this to calculate the longer-term average error. Unfortunately, this soon hits a problem that you can see in the following example.

Period	Demand	Forecast	Error
1	200	0	200
2	400	0	400
3	600	0	600
4	800	2,000	−1,200
Total	2,000	2,000	0
Mean	500	500	0

The positive and negative errors cancel each other, so a very poor forecast can have zero mean error. The mean error does not really show how accurate forecasts are, but it measures bias. If the mean error has a positive value, the forecast is consistently too low: if it has a negative value, the forecast is consistently too high.

We clearly need some other measure of error, and the simplest is to take the absolute values of the errors and calculate a *mean absolute deviation*. This shows how far, on average, a forecast is away from the actual value, so when it equals

6 the forecast is on average 6 away from actual demand. We can also square the errors in each forecast and calculate a *mean squared error*. This value does not have such a clear meaning, but it is useful for other analyses (and we shall meet it again in the next section on linear regression). Whichever measure is used, smaller values mean better forecasts.

Worked example

Hendra Holidays has compared the actual number of holidays it booked each week with its short-term forecasts. What are the errors? What do these errors show?

Week	1	2	3	4	5	6
Demand	101	121	110	98	114	126
Forecast	107	117	112	104	112	120

Solution

The error for the first week is:

$$\text{error} = \text{demand} - \text{forecast} = 101 - 107 = -6$$

The absolute error is 6 and the error squared is 36. The remaining calculations are shown in the spreadsheet of Figure 7.6. From this you can see that the mean error is −0.33, showing that the forecasts are slightly biased, being an average of 0.33 too high. The mean absolute deviation is 4.33, showing that the forecasts are, on average, 4.33 away from actual demand. The mean squared error is 22, but this has no obvious interpretation.

	A	B	C	D	E	F	G	H	I
1	**Errors in forecasts**								
2									
3	**Week**	**1**	**2**	**3**	**4**	**5**	**6**		
4	**Demand**	101	121	110	98	114	126		
5	**Forecast**	107	117	112	104	112	120		
6								**Sums**	**Means**
7	**Error**	**−6**	**4**	**−2**	**−6**	**2**	**6**	**−2**	**−0.33**
8	**Absolute error**	**6**	**4**	**2**	**6**	**2**	**6**	**26**	**4.33**
9	**Error squared**	**36**	**16**	**4**	**36**	**4**	**36**	**132**	**22.00**

Figure 7.6 Calculation of errors

Summary

Inventory control often forecasts time series, which are figures for, say, demand taken at regular periods. These often have underlying patterns, with superimposed

random noise. This noise cannot be forecast and is the reason why forecasts usually contain errors. We can measure these errors using the mean error (which gives bias), mean absolute deviation and mean squared error.

Review questions

7.8 What is a time series?
7.9 Why are forecasts usually wrong?
7.10 Why is the mean error of a forecast of limited use?
7.11 How would you compare different forecasting methods?

Causal forecasting

Causal forecasts look for a cause or relationship that we can use to forecast demand for an item. The demand might, for example, depend on the price being charged. Then we could find the relationship between price and demand, and use this to forecast likely demand at a planned price. This is an example of a true relationship, where changes in price (an independent variable that can take any value) cause changes in demand (a dependent variable that is fixed for any particular value of price). We can illustrate the principles of causal forecasting through linear regression.

Linear regression

Linear regression assumes a dependent variable is linearly related to an independent one, as shown in Figure 7.7. Then it finds the equation of the line of best fit through the data. In particular, it looks for a relationship of the form:

$$\text{dependent variable} = a + b \times \text{independent variable}$$

or

$$y = a + bx$$

where:

x = value of the independent variable
y = value of the dependent variable
a = intercept, where the line crosses the y axis
b = gradient of the line.

Even the best line will not fit the data perfectly, and there is an error at each point. Linear regression finds the values of a and b that minimize some measure

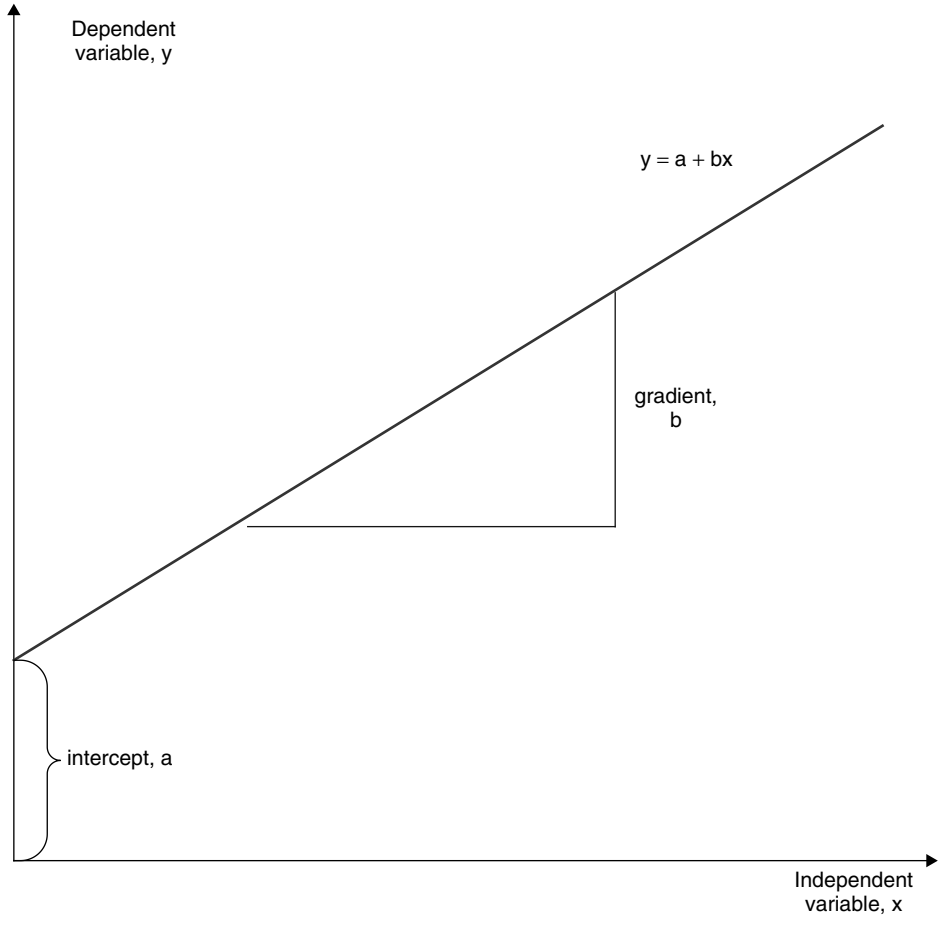


Figure 7.7 Linear relationship used in regression

of the overall error. In practice, the most useful line is the one that minimizes the mean squared error. The equations for this are a bit messy, but we are only interested in the results and, in practice, the calculations are never done by hand.

$$y = a + b \times x$$

$$b = \frac{n \times \sum (x \times y) - \sum x \times \sum y}{n \times \sum x^2 - \left(\sum x\right)^2}$$

$$a = \frac{\sum y}{n} - b \times \frac{\sum x}{n}$$

Worked example

Kurt Steinman's computer supply business is growing, and sales over the past 10 months have been as follows.

Month	0	1	2	3	4	5	6	7	8	9
Demand	3	4	8	10	15	18	20	22	27	28

If Kurt uses linear regression to forecast demand for the next three months, what results will he get?

Solution

The independent variable, x, is the month and the dependent variable, y, is the demand. A lot of software will do the regression calculations, and Figure 7.8 shows a spreadsheet doing the calculations automatically. The data is given in columns A and B. There are some statistics about the results in columns F to I, but you can see the values of a and b in rows 18 and 19. The line of best fit is:

$$y = a + b \times x \text{ or } \qquad \text{Demand} = 2.164 + 2.964 \times \text{month}$$

This equation has been used in column C to forecast values for the first ten months, with the errors shown in column D. The forecasts for the next three

	A	B	C	D	E	F	G	H	I
1	**Linear regression**								
2									
3	**Month**	**Demand**	**Forecast**	**Error**		**Summary output**			
4	0	3	2.16	0.84					
5	1	4	5.13	–1.13		*Regression Statistics*			
6	2	8	8.09	–0.09		Multiple R	0.995	y = 2.164 + 2.964x	
7	3	10	11.05	–1.05		R Square	0.989		
8	4	15	14.02	0.98		Adjusted R Square	0.988		
9	5	18	16.98	1.02		Standard Error	0.993		
10	6	20	19.95	0.05		Observations	10		
11	7	22	22.91	–0.91					
12	8	27	25.87	1.13			*df*	*SS*	*MS*
13	9	28	28.84	–0.84		Regression	1	724.609	724.609
14	10		31.80			Residual	8	7.891	0.986
15	11		34.76			Total	9	732.500	
16	12		37.73						
17							*Coefficients*	*Standard Error*	*t Stat*
18						Intercept, a	2.164	0.584	3.707
19						Gradient, b	2.964	0.109	27.104

Figure 7.8 Linear regression example

months are shown in cells C14:C16 and are calculated from:

$$\text{Month 10:} \quad \text{Demand} = 2.164 + 2.964 \times 10 \quad = 31.8$$

$$\text{Month 11:} \quad \text{Demand} = 2.164 + 2.964 \times 11 \quad = 34.8$$

$$\text{Month 12:} \quad \text{Demand} = 2.164 + 2.964 \times 12 \quad = 37.7$$

If you want to, you can check these calculations, as $n = 10$, $\sum x = 45$, $\sum y = 155$, $\sum x \times y = 942$ and $\sum x^2 = 285$, so that (allowing for rounding):

$$b = (10 \times 942 - 45 \times 155)/(10 \times 285 - 45 \times 45) = 2.96$$

$$a = 155/10 - 2.96 \times 45/10 = 2.18$$

Coefficient of determination

In the last example we found the line of best fit through a set of data. But how well does this line fit the observations? If the errors are small, the line is a good fit to the data; but if the errors are large, even the best line is not very good. To measure the goodness of fit we use the *coefficient of determination*.

The coefficient of determination sees how far the dependent values are away from their mean. Some of the variation from the mean is explained by the linear relationship – some is unexplained and is due to random noise. The coefficient of determination gives the proportion of the total error that is explained by the linear relationship. It has a value between zero and one. If it is near to 1, most of the variation is explained by the regression, there is little noise and the straight line is a good fit to the data. If the value is near to zero, most of the variation is unexplained, there is a lot of random noise, and the line is not a good fit. (Actually, because of the problem with positive and negative errors cancelling the coefficient works with the squared errors, but this is only a detail.)

The coefficient of determination is often called 'r^2', and if you look in Figure 7.8 you can see 'R Square 0. 989' in cells F7:G7. This shows that 98.9 per cent of the variation of the dependent variable from its mean is explained by the linear relationship. Only 1.2 per cent of the variation is due to noise, so the errors are very small. You can check these results in column H where the entries 'SS' show the sum of squares: the total of the squared variation from the mean is 732.5, of this 724.609 is explained by the regression leaving a residual of 7.891 that comes from the error. So the coefficient of determination is $724.609/732.5 = 0.989$. In general, values of r^2 above, say, 0.5 show good results with reasonably small errors.

Coefficient of correlation

A second useful measure in regression is the *coefficient of correlation* which asks the question 'Are x and y linearly related?' The coefficients of correlation and

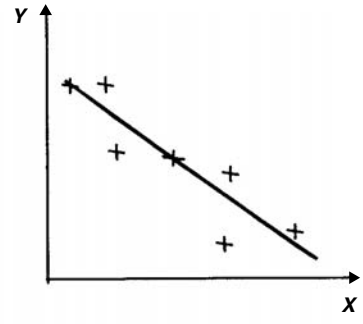
determination answer very similar questions, and it is easy to show that:

$$\text{coefficient of determination} = (\text{coefficient of correlation})^2$$

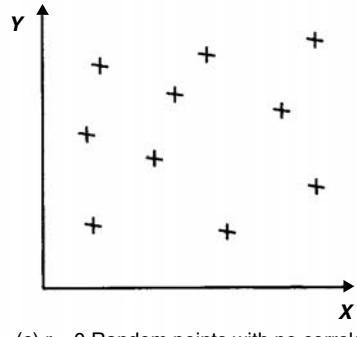
As the coefficient of determination is r^2, the coefficient of correlation is r, and has a value between +1 and −1 (illustrated in Figure 7.9):



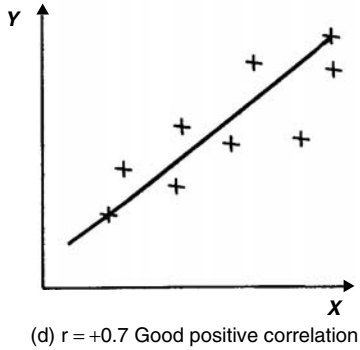
(a) r = –1 Perfect negative correlation



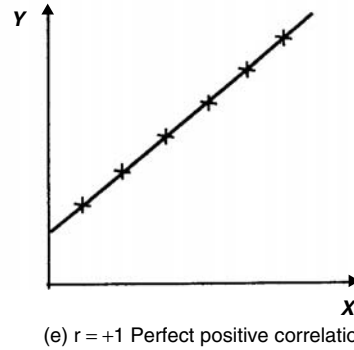
(b) r = –0.7 Good negative correlation



(c) r = 0 Random points with no correlation



(d) r = +0.7 Good positive correlation



(e) r = +1 Perfect positive correlation

Figure 7.9 Coefficient of correlation

- A value of r = 1 shows the two variables have a perfect linear relationship with no noise at all, and as one increases so does the other.
- A low positive value of r shows a weak linear relationship.
- A value of r = 0 shows there is no correlation at all between the two variables and no linear relationship.
- A low negative value of r shows a weak linear relationship.
- A value of r = −1 shows the two variables have a perfect linear relationship and as one increases the other decreases.

A correlation coefficient near to +1 or −1 shows a strong linear relationship. However, when r is between 0.7 and −0.7 the coefficient of determination is less than 0.49 and less than half the variation is explained by the regression. As a rule of thumb, linear regression is not very reliable when the coefficient of correlation is between about 0.7 and −0.7. In Figure 7.8, cells F6:G6 show that 'Multiple R 0.995'. This is the coefficient of correlation, which again confirms a very good fit, with the demand rising with month.

Worked example

Over the past 16 weeks Burridge Transport Ltd has recorded the following number of loads moved for a particular customer. What can the company learn from these figures?

Week	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Demand	194	187	163	172	160	154	143	156	131	167	143	135	126	95	112	107

Solution

There seems to be a fairly steady decline in work for this customer, but we can see this more clearly from the linear regression in Figure 7.10. This shows the line of best fit through the data as:

$$\text{Demand} = 191.875 - 5.331 \times \text{week}$$

So demand is falling by an average of 5.331 units a week. Using this equation to forecast the volume of business over the next few weeks gives 101.25, 95.92, 90.59, 85.26, 79.93 etc. (shown in cells C20:C24). The coefficient of correlation is −0.905 which shows the line is a very good fit to the data; the coefficient of determination is 0.820, which shows that 82 per cent of the variation from the mean is explained by the regression. You can see these results more clearly in the graphs shown in Figure 7.11.

	A	B	C	D	E	F	G	H	I
1	**Linear regression**								
2									
3	**Week**	**Demand**	**Forecast**	**Error**		**Summary output**			
4	1	194	186.54	7.46					
5	2	187	181.21	5.79		*Regression Statistics*			
6	3	163	175.88	–12.88		Multiple R	–0.905	y = 191.875-5.331x	
7	4	172	170.55	1.45		R Square	0.820		
8	5	160	165.22	–5.22		Adjusted R Square	0.807		
9	6	154	159.89	–5.89		Standard Error	12.322		
10	7	143	154.56	–11.56		Observations	16		
11	8	156	149.23	6.77					
12	9	131	143.90	–12.90			*df*	*SS*	*MS*
13	10	167	138.57	28.43		Regression	1	9662.224	9662.224
14	11	143	133.24	9.76		Residual	14	2125.713	151.837
15	12	135	127.90	710		Total	15	11787.938	
16	13	126	122.57	3.43					
17	14	95	117.24	–22.24			*Coefficients*	*Standard Error*	*t Stat*
18	15	112	111.91	0.09		Intercept, a	191.875	6.462	29.694
19	16	107	106.58	0.42		Gradient, b	–5.331	0.668	–7.977
20	17		101.25						
21	18		95.92						
22	19		90.59						
23	20		85.26						
24	21		79.93						

Figure 7.10 Linear regression for Burridge Transport Ltd

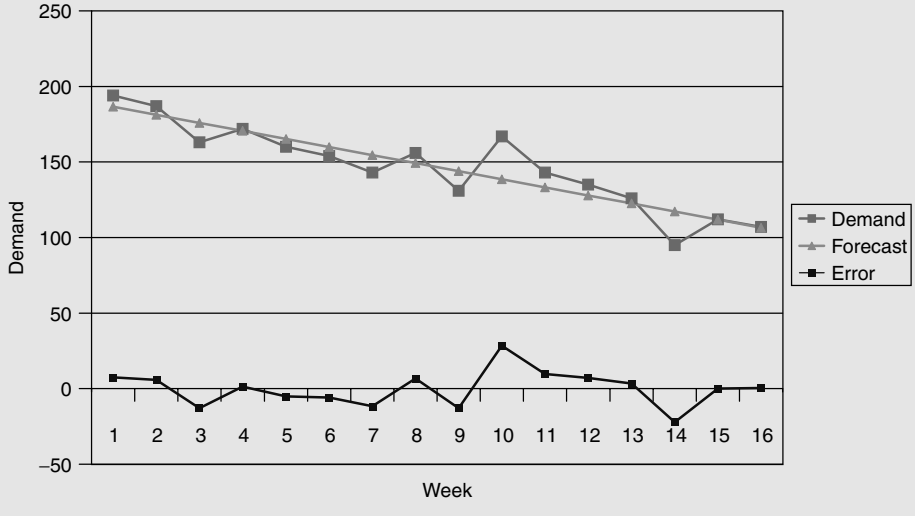


Figure 7.11 Results for Burridge Transport Ltd

Extensions to linear regression

There are several extensions to the basic linear regression model. One considers multiple linear regression (invariably called *multiple regression*), which looks

for a linear relationship between a dependent variable and several independent ones:

$$y = a + b_1 \times \text{variable } 1 + b_2 \times \text{variable } 2 + b_3 \times \text{variable } 3 + b_4 \times \text{variable } 4 \ldots$$

The sales of a product, for example, might depend on its price, the advertising budget, number of suppliers, local unemployment rate, and so on.

Another extension looks at non-linear regression, where a more complicated line is fitted to data. A population might, for example, grow exponentially, with:

$$\text{population, } y = a \times b^{\text{period}}$$

where a and b are constants. Spreadsheets and other packages have a range of standard procedures for dealing with more complicated regressions of this type.

Summary

Causal forecasts look for relationships between variables. We illustrated this by linear regression, which finds the line of best fit through a set of data, and allows a dependent variable to be forecast from a related independent variable. The coefficient of determination shows how well a regression line fits the data, while the coefficient of correlation shows how strong a linear relationship is.

Review questions

7.12 How does causal forecasting work?

7.13 What is 'linear regression'?

7.14 What is measured by the coefficient of determination?

7.15 The coefficient of correlation sees if changes in the independent variable cause changes in the dependent variable. Do you think this is true?

Projective forecasting

Projective forecasting uses historical values of demand to forecast future demand – projecting past patterns into the future. We will describe the four most widely used methods based on a simple average, moving average, exponential smoothing and a model for seasonality and trend.

Simple average

If you want to know the demand for a product in the future, an obvious method is to find the average demand in the past. This is the approach of forecasting using a simple average.

Worked example

Use simple averages to forecast demand for period 6 of the following two time series. How accurate are the forecasts? What are the forecasts for period 27?

Period	1	2	3	4	5
Series 1	49	50	49	52	50
Series 2	70	33	76	29	42

Solution

For the first series the average demand is:

$$(49 + 50 + 49 + 52 + 50)/5 = 250/5 = 50$$

This is the forecast demand for period 6. Similarly, the total demand for the second series is also 250, so the forecast for period 6 is again 50. Although the forecasts are the same, there is clearly less noise in the first series than the second. We should, therefore, be more confident in the result for the first series and expect a smaller error.

Simple averages assume the demand is constant over time, so the forecasts for period 27 are the same as the forecasts for period 6, at 50.

Simple averages are easy and can give good results when demand is stable. Unfortunately, it does not work well if the demand pattern changes. Then older data tends to swamp the latest figures and the forecast does not respond quickly enough to changes. Suppose, for example, demand for an item has been constant at 50 units a week for the past year. A simple average would forecast demand for week 53 as 50 units. If the demand in week 53 suddenly rises to 100 units, a simple average gives a forecast for week 54 of:

$$(52 \times 50 + 100)/53 = 50.94$$

Doubling of demand has increased the forecast by less than 2 per cent. If demand continues at 100 units a week, subsequent forecasts are 51.85, 52.73, 53.57, 54.39, 55.17, and so on. The forecasts are clearly rising, but the response is very slow.

Simple averages can only really be used for demand that is stable over long periods. In practice, very few demands are this stable, so the method has limited value.

Moving average

The problem with a simple average is that old data – which is probably out of date – swamps newer, more relevant data. One way around this is to ignore old data and only use the most recent values in forecasts. We might use the average

demand over, say, the past 12 months for our forecast and ignore any older data. This is the basis of a moving average.

Worked example

Demand for an item over the past 6 months has been as follows:

Month	1	2	3	4	5	6	7	8
Demand	203	194	188	206	173	119	209	194

The market for this item is unstable, and any data over 3 months old is unreliable. Use a moving average to forecast demand for the item.

Solution

Only data more recent than three periods is relevant, so we can use a three-period moving average for the forecast. If we consider the situation at the end of month 3, the forecast for month 4 is:

$$\text{Forecast} = (203 + 194 + 188)/3 = 195$$

At the end of month 4 we know that actual demand is 206, and we can update the forecast for month 5 to give:

$$\text{Forecast} = (194 + 188 + 206)/3 = 196$$

Then at the end of month 5 we can update the forecast for month 6 to:

$$\text{Forecast} = (188 + 206 + 173)/3 = 189$$

The results are summarized in the following table.

Month	1	2	3	4	5	6	7	8	9
Demand	203	194	188	206	173	119	209	194	–
Forecast	–	–	–	195	196	189	166	167	174

You can see from this worked example how a moving average forecast responds to changing demand, with a high demand moving the forecast upwards and a low demand moving it downwards. At the same time the forecast is smoothing out variations, so that it does not blindly follow the effects of random noise. The rate at which a forecast responds to changing demand is called its *sensitivity*. With a moving average we can adjust this by altering the number of periods we average. Taking the average of a small number of observations gives a responsive forecast that quickly follows changes in demand but it might be too sensitive and follow atypical values and random fluctuations. Taking the average of more observations gives a less sensitive forecast that smoothes out random variations but it may not follow genuine changes in demand. We need a compromise that gives reasonable results, and a typical value is around six periods.

Worked example

Demand for an item over the past 11 weeks is as follows. Use moving averages over different periods to find one week ahead forecasts.

Week	1	2	3	4	5	6	7	8	9	10	11
Demand	42	33	36	45	54	63	69	72	75	78	98

Solution

If we start with a three-period moving average, the first forecast we can find is for week 4, which is $(42 + 33 + 36)/3 = 37$. The results for three, four and six-period moving averages are given in Figure 7.12. As you can see from the

Moving averages

Week	Demand	Three period	Four period	Six period
1	42			
2	33			
3	36	37.00		
4	45	38.00	39.00	
5	54	45.00	42.00	
6	63	54.00	49.50	45.50
7	69	62.00	57.75	50.00
8	72	68.00	64.50	56.50
9	75	72.00	69.75	63.00
10	78	75.00	73.50	68.50
11	98	83.67	80.75	75.83

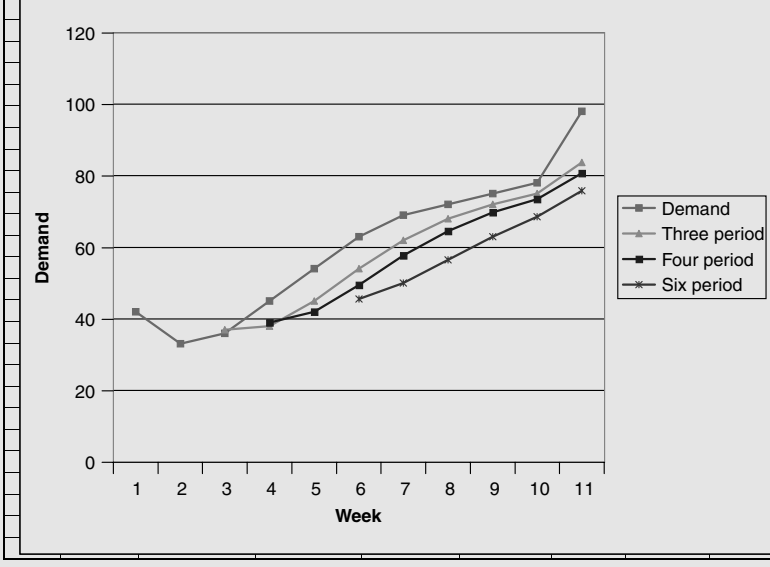


Figure 7.12 Comparison of moving averages

graphs, there is an upward trend in demand. The three-period moving average is most responsive to change and is following this trend fastest; the six-period moving average is least responsive.

Exponential smoothing

Moving averages overcome some of the problems with simple averages, but they still have two major defects:

- all historical values are given the same weight;
- the method only works well with relatively constant demand.

Both of these problems are avoided by *exponential smoothing*. This is based on the idea that as data gets older it becomes less relevant and should be given less weight. So the method gives high weight to the most recent data, and this weight declines exponentially with age, as shown in Figure 7.13.

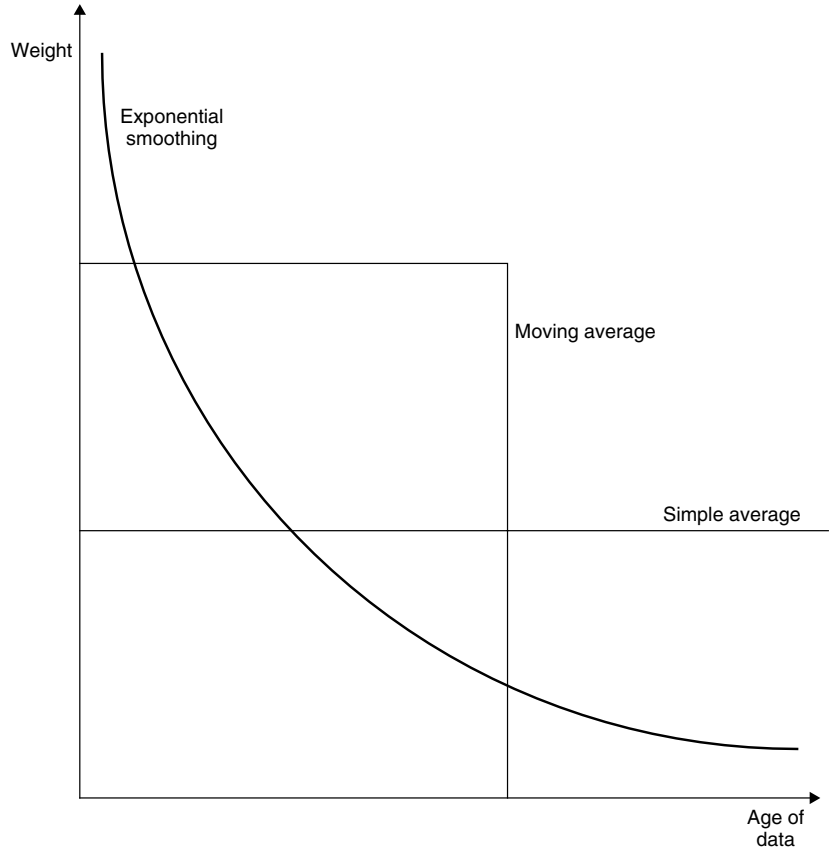


Figure 7.13 Varying weight given to data with age

In practice, we can get this declining weight using only the latest demand figure and the previous forecast. To be precise, we take a proportion, α, of the latest demand and add a proportion, $1 - \alpha$, of the previous forecast:

$$\boxed{\text{new forecast} = \alpha \times \text{latest demand} + (1 - \alpha) \times \text{previous forecast}}$$

In this equation, α is the smoothing constant that is usually given a value between 0.1 and 0.2.

You can see how exponential smoothing adapts to changes in demand with a simple example. Suppose a forecast was optimistic and suggested a value of 100 for a demand that actually turns out to be 90. Taking a value of $\alpha = 0.2$, the forecast for the next period is:

$$\begin{aligned}\text{new forecast} &= \alpha \times \text{latest demand} + (1 - \alpha) \times \text{previous forecast}\\ &= 0.2 \times 90 + (1 - 0.2) \times 100\\ &= 98\end{aligned}$$

The lower demand has clearly reduced the forecast for the next period.

Worked example

An item has the following weekly demand. Use exponential smoothing with $\alpha = 0.2$ and an initial forecast for week 1 of 102 units to find one period ahead forecasts.

Week	1	2	3	4	5	6	7	8
Demand	107	115	94	89	98	91	101	112

Solution

Starting at the end of week 1, we know that the previous forecast was 102 and the latest demand is 107. We can use these to find the new forecast for week 2 as:

$$\text{New forecast} = 0.2 \times 107 + 0.8 \times 102 = 103$$

Then at the end of week 2 when we know the demand was 115, we can find a new forecast for week 3 as:

$$\text{New forecast} = 0.2 \times 115 + 0.8 \times 103 = 105.4$$

These forecasts are summarized in the following table:

Week	1	2	3	4	5	6	7	8	9
Demand	107	115	94	89	98	91	101	112	
Forecast	102	103	105.4	103.1	100.3	99.8	98.0	98.7	101.3

The value given to the smoothing constant sets the sensitivity of the forecast. A high value of α (say 0.3 to 0.35) puts more weight on the latest demand and gives a more responsive forecast: a lower value (say 0.1 to 0.15) puts more weight on the old forecast and gives a less responsive forecast. Again, we need to compromise between a responsive forecast that might follow random fluctuations, and an unresponsive one that might not follow real patterns.

Worked example

The following time series has a clear step upwards in demand in month 4. Use an initial forecast of 50 to compare exponential smoothing forecasts with varying values of α.

Month	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Demand	48	50	150	145	155	150	148	152	150	149	150	161	155	148

Solution

Figure 7.14 shows the results for various values of α. As you can see from the graphs, all the forecasts eventually follow the sharp step and raise forecasts to around 150. Higher values of α make this adjustment more quickly and give a more responsive forecast.

Although higher values of α give more responsive forecasts, they do not necessarily give more accurate ones, as they follow the random effects of noise. One way of setting a reasonable value for α is to test several values over a trial period and choose the one that gives smallest errors. An alternative is to vary the value depending on circumstances. If, for example, forecast errors begin to rise, the forecast may be adapting to a real change in demand and a more sensitive forecast would give better results. There are several ways of keeping a check on errors – including the measures of error that we used previously. A simpler approach uses a tracking signal. Many of these have been suggested, with a common one defined as:

$$\text{Tracking signal} = \frac{\text{sum of forecast errors}}{\text{mean absolute deviation}}$$

	A	B	C	D	E	F	G	H	I
1	**Exponential smoothing**								
2									
3			**Forecast with alpha equal to**						
4	**Period**	**Demand**	**0.1**	**0.2**	**0.3**	**0.4**			
5	1	48	50.00	50.00	50.00	50.00			
6	2	50	49.80	49.60	49.40	49.20			
7	3	150	49.82	49.68	49.58	49.52			
8	4	145	59.84	69.74	79.71	89.71			
9	5	155	68.35	84.80	99.29	111.83			
10	6	150	77.02	98.84	116.01	129.10			
11	7	148	84.32	109.07	126.20	137.46			
12	8	152	90.69	116.86	132.74	141.67			
13	9	150	96.82	123.88	138.52	145.80			
14	10	149	102.14	129.11	141.96	147.48			
15	11	150	106.82	133.09	144.07	148.09			
16	12	161	111.14	136.47	145.85	148.85			
17	13	155	116.13	141.37	150.40	153.71			
18	14	148	120.01	144.10	151.78	154.23			
19	15		122.81	144.88	150.64	151.74			

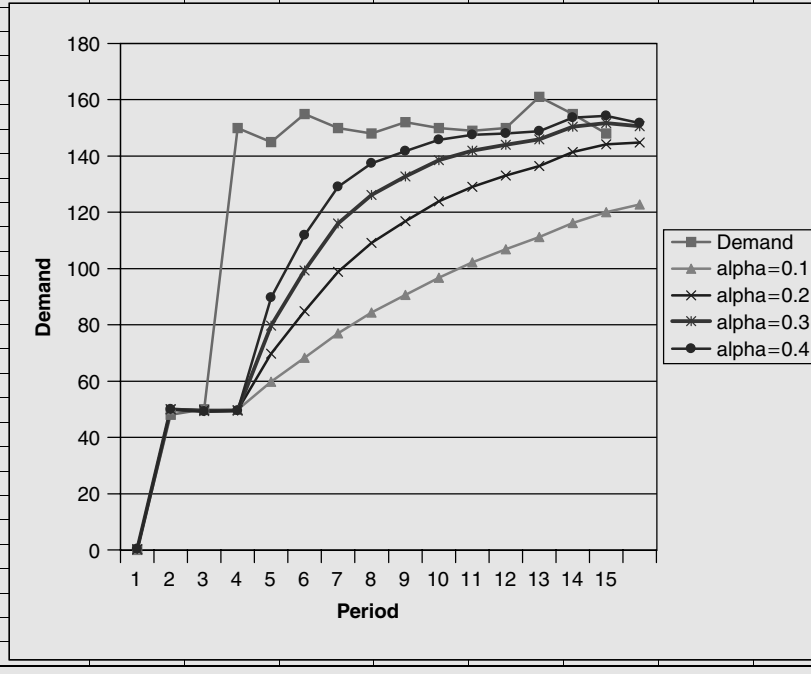


Figure 7.14 Effect of varying the smoothing constant

When the forecast error remains small, this tracking signal has a value close to zero. But if the error get bigger, the value of the tracking signal increases. When it reaches some predefined limit around, say, 2.5 remedial action is needed, perhaps increasing the value of α.

Models for seasonality and trend

The projective methods we have described so far only work well with constant demand, and they need some adjustments to deal with other patterns. In this section we describe a model to use when there is seasonality and trend. Here 'trend' is the amount that demand grows between two consecutive periods. If two consecutive periods have demands of 50 and 60 the trend is 10: if two consecutive periods have demands of 110 and 100 the trend is −10. 'Seasonality' is a regular cyclical pattern. Each cycle repeats the same general pattern, so that sales of a newspaper follow a weekly cycle with highest sales at the weekend. Each cycle contains a number of 'seasons', and we can measure the variation in demand using seasonal indices.

$$\text{Seasonal Index} = \frac{\text{Seasonal value}}{\text{Deseasonalized value}}$$

Suppose a newspaper has average daily sales of 100 copies in a particular area, but this rises to 200 copies on Saturday and falls to 50 copies on Monday and Tuesday. The deseasonalized value is 100, the seasonal index for Saturday is $200/100 = 2.0$, the seasonal indices for Monday and Tuesday are $50/100 = 0.5$, and seasonal indices for other days are $100/100 = 1.0$.

The easiest way of forecasting complex time series is to split demand into separate components, and forecast each component separately. Then we recombine the separate components to get the final forecast. To be specific, we shall split the demand into four components (shown in Figure 7.15):

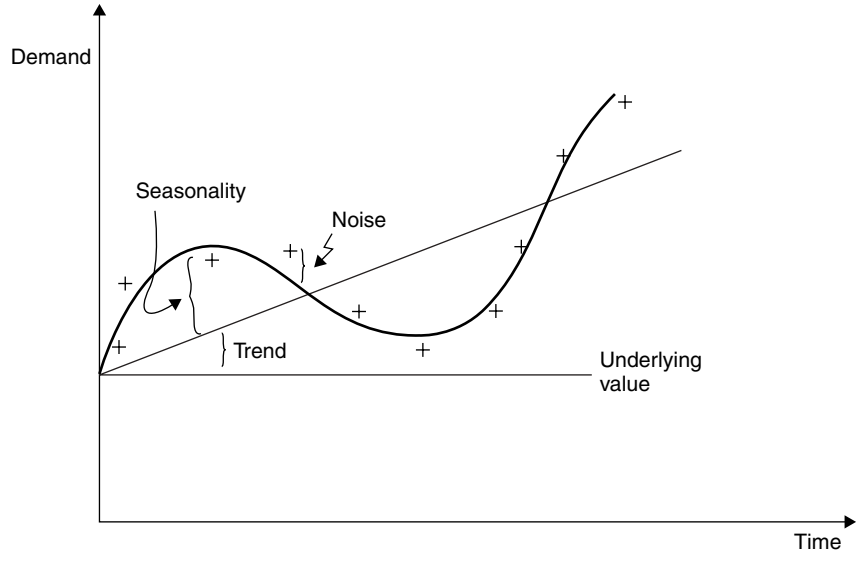


Figure 7.15 Components of a time series with seasonality and trend

- *Underlying value* is the basic demand that must be adjusted for seasonality and trend.
- *Trend* shows the change in demand between periods.
- *Seasonality* is the cyclical variation around the trend.
- *Noise* is the random noise whose effects we cannot explain.

Then demand is made up of an underlying value, with added trend, multiplied by a seasonal index, and added noise.

$$\text{Demand} = (\text{underlying value} + \text{trend}) \times \text{seasonal index} + \text{noise}$$

There are several ways of doing the calculations for this kind of forecast, but a straightforward one has the following seven steps:

1. Use linear regression on the time series to find the underlying value (the intercept, a) and the trend (the gradient, b).
2. Use this regression to find a deseasonalized value for each period.
3. We now have the original seasonal data for each period, and corresponding deseasonalized values from regression. Dividing the first of these by the second gives a seasonal index for each period.
4. Look at the data to see how many seasons there are in each cycle.
5. Find the average seasonal index for each season of the cycle, i.e. the average index for the first season, second season, etc.
6. Find deseasonalized values for the future by projecting the regression line.
7. Calculate the final forecast by multiplying these deseasonalized values by the corresponding seasonal indices.

Worked example

Over the past 12 quarters the demand for an item has been as follows:

Quarter	1	2	3	4	5	6	7	8	9	10	11	12
Demand	602	620	304	396	798	804	602	630	941	896	664	736

How would you forecast demand for the next five quarters?

Solution

For this forecast we can use the standard approach described above, and these calculations are shown in the spreadsheet of Figure 7.16. The historical data is in rows 4 to 15, and the forecasts are in rows 16 to 20.

	A	B	C	D	E	F	G
1	**Forecasting with seasonality and trend**						
2							
3	**Quarter**	**Season in cycle**	**Demand**	**Deseasonalized value**	**Seasonal index**	**Forecast**	
4	1	1	602	511.60	1.18		
5	2	2	620	539.69	1.15		
6	3	3	304	567.78	0.54		
7	4	4	396	595.86	0.66		
8	5	1	798	623.95	1.28		
9	6	2	804	652.04	1.23		
10	7	3	602	680.13	0.89		
11	8	4	630	708.21	0.89		
12	9	1	941	736.30	1.28		
13	10	2	896	764.39	1.17		
14	11	3	664	792.48	0.84		
15	12	4	736	820.56	0.90		
16	13	1		848.65	1.24	1056.188	
17	14	2		876.74	1.18	1038.655	
18	15	3		904.83	0.75	681.162	
19	16	4		932.91	0.82	762.217	
20	17	1		961.00	1.24	1196.013	
21							
22							
23	*Regression Statistics*				*Coefficients*		
24	Multiple R	0.541002		Intercept,a	483.515152		
25	R Square	0.292683		Gradient, b	28.0874126		
26	Adjusted R Square	0.221951					
27	Standard Error	165.1157					
28	Observations	12					
29							
30							
31							
32							
33							
34							
35							
36							
37							
38							
39							
40							
41							
42							
43							
44							
45							
46							
47							
48							
49							
50							
51							
52							

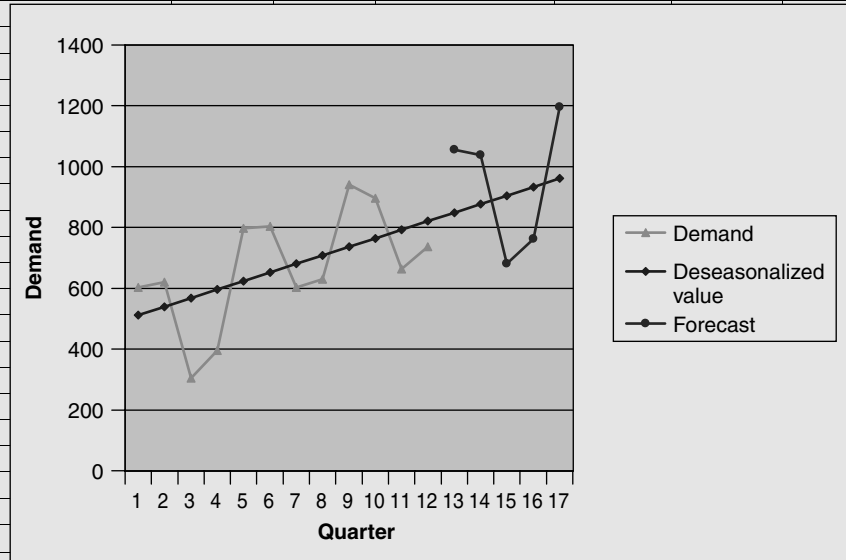


Figure 7.16 Forecasting with seasonality and trend

1. Use linear regression on the time series to find the underlying value (the intercept, a) and the trend (the gradient, b). The results of the regression are in rows 23 to 28 which show that the underlying value is 483.52 and the trend is 28.09, and:

$$\text{Demand} = 483.52 + 28.09 \times \text{quarter}$$

   The coefficient of correlation is low, but we would expect this with strongly seasonal data.

2. Use this regression to find a deseasonalized value for each period. These are shown in column D for quarters 1 to 17.
3. Dividing the actual demand by the deseasonalized value gives a seasonal index for each of the first 12 quarters, which are shown in column E.
4. Look at the data to see how many seasons there are in each cycle. You can see from the graph of demand, and also the seasonal indices, that there are four quarters in each cycle. So we have data for three annual cycles.
5. Find the average seasonal index for each season of the cycle. The first seasons in the three cycles are quarters 1, 5 and 9, and the average seasonal index for these three is $(1.18 + 1.28 + 1.28)/3 = 1.24$. Average seasonal indices for the other three quarters are found in the same way, as 1.18, 0.75 and 0.82 (shown in cells E16:E20).
6. Find deseasonalized values for the future by projecting the regression line. The deseasonalized values for quarters 13 to 17 are shown in cells D16:D20.
7. Calculate the final forecast by multiplying these deseasonalized values by the corresponding seasonal indices. The forecasts are in rows 16 to 20. For quarter 13, for example, the deseasonalized value is 848.65, the seasonal index for the first quarter is 1.24, and multiplying these together gives a forecast of 1056.19 (allowing for rounding). Forecasts for the next four quarters are shown in cells F17:F20.

Summary

Projective forecasting extends past patterns of demand into the future. The easiest approach uses a simple average, but this can be unresponsive to change and only gives good results for stable demand. A moving average takes the average demand over the latest periods as a forecast. Altering the number of periods that are averaged changes the sensitivity. Exponential smoothing gives a reducing weight to data as its age increases, and the sensitivity can be adjusted by changing the value of the smoothing constant. Time series with seasonality and trend can be split into four components of underlying value, trend, seasonal index and noise. The first three of these can be forecast separately and combined to give a final forecast.

Review questions

7.16 Why is a simple average of limited use for forecasting?
7.17 How can you make a moving average forecast more responsive?
7.18 Why is the forecasting method called 'exponential smoothing'?
7.19 How can you make exponential smoothing more responsive?
7.20 Why is ordinary exponential smoothing not used for demand figures with seasonality and trend?
7.21 How can you forecast demand that has trend but no seasonality?

Planning forecasts

We have outlined some useful methods of forecasting demand, but many things may still go wrong. You can see a lot of evidence for this with poor weather forecasts, projections for national economic growth, planned company profits, sales projections, winners of football matches, and so on. But what exactly goes wrong? Why do economists who use huge computer systems with virtually limitless information and the best models available still make mistakes when forecasting next quarter's gross domestic product? There are, unfortunately, many answers to this. The obvious one is the random noise that we cannot explain, but there are other factors that are generally based on errors in the forecasting model, errors in the data, or errors in interpreting results.

- The noise is a completely random effect that is caused by many factors, such as varying customer demand, hours worked, speed of working, weather, rejections at inspections, time of year, and so on. Although we will try to explain as much of the variation in demand as possible, the noise is, by definition, random and inexplicable.
- Errors in the model can include using the wrong type of model, not recognizing real patterns in demand, analysing patterns where none exists, unnecessarily complex models, ignoring external influences like the business cycle, ignoring internal influences like product life cycles or promotions, being too inflexible, using the wrong parameters, presenting results in the wrong way, taking too long to get results, etc.
- Errors in the data can include simple mistakes, missing data, delays in getting up-to-date results, poor communications and errors in transmitting data, over-reliance on one source of data, conflicting data from different sources, inconsistent results when the method of collection or analysis changes, not collecting some relevant data, measuring data at the wrong point in the supply chain, making invalid assumptions, using inappropriate analyses or conventions, etc.
- Errors in interpreting results can include over-confidence in a poor forecasting method, over-confidence in flawed data, missing real changes to customer demands and opinions, not recognizing the impact of external factors, ignoring

special conditions, mistaking precision for accuracy, not recognizing real patterns in the forecasts, imagining patterns where none exists, not including lost sales, not taking account of likely errors, etc.

Such errors can bring major problems with, most obviously, unrealistically high forecasts of demand for an expensive item that remains unsold, or unrealistically low forecasts of demand for essential items. These are the two main consequences of poor forecasts. If the forecast is too high, there is too much stock with high holding costs and a greater chance of obsolescence and loss; if the forecast is too low, the stock is too low and there are high shortage costs from lost sales and dissatisfied customers. When there is uncertainty, managers may wonder whether it is better to over-estimate demand or under-estimate it. Bearing in mind the high costs of shortage, we might be tempted to over-estimate and build-up more safety stock. There is an argument, however, that this gives unnecessarily high stock. Low forecasts will usually only lead to relatively small shortages, and the best way of dealing with these is equitably to share out the available stock between orders. Then each customer is largely satisfied, and will hopefully wait for the remainder in a back-order (see Thacker, 2001b).

Of course, the best option is not to make mistakes in the forecasts, and we can give some guidance on how to avoid the most common mistakes. Many of these points are obvious, but people still forget them or perhaps imagine that they do not have the time to consider them.

The starting point for any forecast is clearly to define the purpose of the forecast. What are we trying to forecast? Why? How will we use the results? When will we need them? If you do not have a clear idea of the results you want, or the reason that you want them, you cannot hope to do appropriate analyses. There is always a temptation to look at some easily available data, put it into a spreadsheet and 'do a bit of forecasting'. The results are probably useless, but any numbers that appear on a computer screen seem somehow persuasive.

There is also a temptation to see a forecast as an end in itself – once we have a set of numbers the job is finished. But the whole purpose of forecasting is not to produce numbers, but to give some information that can help managers make informed decisions. This means that we should involve other people, particularly those who have most knowledge and experience about the area, and those who will use the results. Remember that different people want different types of result, and while some will only want a single overall figure, others will want details of the methods, data, assumptions and potential problems. Remember also the use of forecasts, with sales people, for example, tending to prefer optimistic forecasts, while finance people tend to prefer pessimistic ones. People who are paid a bonus for exceeding their forecast performance will always look for very low forecasts.

Another important point is the accuracy needed in forecasts. Aggregate forecasts are more accurate than those for specific items, and paint manufacturers, for example, can get more accurate forecasts of their total sales of paint than for the sales of each particular colour. The time horizon covered also has an effect, with short-term forecasts being more accurate than long-term ones. On the whole – but not inevitably – more accurate forecasts need more effort, cost more and take

longer. So we have to balance the costs and benefits: there is no point in spending a lot of time and money on an unimportant forecast, but some forecasts are very important and we should put in enough effort to get the best possible results.

This question of accuracy also affects expectations. Some people can develop too high expectations of a forecast and are disappointed when there are any of the inevitable errors. We should anticipate some errors, try to estimate the likely effects, make contingency plans, avoid making firm commitments based on forecasts, warn people that the results can never be perfect, and always monitor forecast performance.

This brings us to the actual choice of a forecasting method. If there is historical data, quantitative methods are much more reliable. But there are many of these available, and the final choice depends on factors like the time covered, data available, variability of demand, accuracy needed, the amount of time and money available, expertise, software used, current practices, etc. Do not assume that a good or complicated method inevitably gives good results and remember that you can often get reasonable results from very simple procedures. Before making a final decision, you should test the likely choice over a typical historical period to make sure that it gives good results – and results that satisfy all the requirements. This is also the time to collect any missing data and test different parameter values.

Now that we have done all the planning and made all the decisions, it is time to implement our forecasting method. But we have to do this carefully, perhaps running the new methods and old ones in parallel for some time and making sure that the new methods are working properly and really do give an improvement. After this, there is the continual monitoring of performance, checking for errors, and continual search for improvements and better parameter values.

Summary

Forecasts inevitably contain errors, which might be caused by noise, errors in models, errors in data or errors in interpreting results. To minimize the effects of these errors, we need to plan the forecasts. This includes decisions about the purpose of the forecast, people involved, results needed, accuracy, available data, etc. The aim of forecasting is not to produce a series of figures, but to get information that can help managers make informed decisions.

Review questions

7.22 Is it generally better to over-estimate or under-estimate demand?

7.23 The most important thing for successful forecasting is choosing the right method. Do you think this is true?

Chapter review

- All decisions about stocks become effective at some point in the future, so they rely on forecasts of prevailing conditions. There are many methods of

forecasting, but none is always the best. We can classify the methods in several ways, with useful ones describing the time they look ahead (in the short, medium or long term), and the overall approach (judgement, causal or projective).

- Judgemental forecasts rely on subjective views and opinions of experts. They are the only possible methods when there is no reliable historical data. The most common methods are personal insight, panel consensus, market surveys, historical analogy and the Delphi method.
- Inventory control often forecasts time series, which are figures for, say, demand taken at regular periods. These generally have underlying patterns, with superimposed random noise.
- The noise cannot be forecast and is the reason why forecasts usually contain errors. We can measure these errors using the mean error, mean absolute deviation and mean squared error.
- Causal forecasts look for relationships between variables. We illustrated this by linear regression, which finds the line of best fit through a set of data. We can use this relationship to forecast the value of a dependent variable from a related independent variable. The coefficients of determination and correlation show how good the result is likely to be.
- Projective forecasting extends past patterns of demand into the future. The easiest approach uses a simple average, but this can be unresponsive to change and only gives good results for stable demand.
- Other methods include a moving average, which finds the average demand over a specified number of periods of the latest demands, and exponential smoothing, which gives a reducing weight to data as it gets older.
- The easiest way of forecasting time series with seasonality and trend divides the values into four components of underlying value, trend, seasonal index and noise. We forecast the first three of these separately and combine them to give overall forecasts.
- Forecasts inevitably contain errors, which might be caused by noise, errors in models, errors in data or errors in interpreting results. To minimize the effects of these errors, we need to plan the forecasting. This includes decisions about the purpose of the forecast, people involved, results needed, accuracy, available data, etc.

Project

You can find long time series published by the government, or other organizations such as the Stock Exchange, the United Nations or the World Bank. The aim of this project is to see how well you can forecast these values. You might start by seeing how share prices – either for specific companies or aggregate

figures – have moved over the past few years. How well do you think you can forecast them in the future? Now compare these results with, say, government figures for unemployment, GDP, trade, travel, sales of commodities, or some other time series. Can you find patterns to forecast future values? What problems and successes do you have?

Problems

7.1 Gerry Hunter and Sonya Harding each forecast the monthly sales of their main product for a year in advance. At the end of the year they compared their results, as shown below. Who do you think produced the better forecasts?

Month	1	2	3	4	5	6	7	8	9	10	11	12
Demand	40	44	52	38	28	30	27	38	54	45	51	56
Gerry	34	46	48	44	34	32	30	32	36	38	41	43
Sonya	30	40	44	48	38	36	42	40	42	44	44	44

7.2 Pentishaw Publications is producing a limited edition print and wants to know how much to spend on advertising. Its aim is to attract as many customers as possible, up to the print run of 400 copies. For the past ten prints the spending on advertising (in thousands of pounds) and resulting sales are shown in the following table. How much would you spend on advertising?

Spending	6	9	2	12	4	9	7	5	10	12
Sales	275	350	150	525	225	400	325	300	400	500

7.3 In the special products section of Kleinmann Industries, production of an item depends on the number of shifts worked in a month. Data for the past 9 months is shown below. If 400 units are needed next month, how many shifts should Kleinmann plan? How reliable is this result?

Month	1	2	3	4	5	6	7	8	9
Shifts worked	50	70	25	55	20	60	40	25	35
Units made	352	555	207	508	48	498	310	153	264

7.4 Martin Chua has collected the following data relating to sales. What information can he get from this?

Month	1	2	3	4	5	6	7	8	9	10
Sales	220	360	460	370	260	320	330	380	290	340
Market	4	8	9	2	1	5	2	9	5	6
Personnel	21	14	56	53	13	40	36	20	19	24
Competition	109	132	150	165	185	95	130	145	95	110

7.5 The following table shows monthly demand for a product over the past year. Use 2, 3 and 4 period moving averages to produce one period ahead forecasts, and say which gives the best results:

Month	1	2	3	4	5	6	7	8	9	10	11	12
Demand	32	28	24	30	36	42	46	48	50	52	64	76

7.6 See if you can use a moving average to deseasonalize the following time series:

Period	1	2	3	4	5	6	7	8	9	10
Demand	188	75	130	220	80	133	225	75	140	240

7.7 Use exponential smoothing with α in the range 0.1 to 0.4 and an initial value of 69 to produce one period ahead forecasts for the following time series. Which result is best?

Period	1	2	3	4	5	6	7	8	9	10	11
Demand	71	78	101	122	94	89	73	69	59	64	68

7.8 Anna Gupta has recorded quarterly demand for a product over the past three years as follows. How could she forecast demand over the next two years?

Quarter	1	2	3	4	5	6	7	8	9	10	11	12
Demand	190	143	101	168	228	184	138	213	270	210	150	255

7.9 Over the past two years Lars Suntersonn had the following demands for a product. What forecasts would you give for demand over the following two years?

Month	Jan	Feb	Mar	Apr	May	June	July	Aug	Sept	Oct	Nov	Dec
Year 1	200	186	178	174	189	207	215	233	220	218	220	243
Year 2	227	203	211	190	222	236	230	253	257	239	247	282

Discussion questions

7.1 Can you find some real examples of poor forecasts? How have these affected the organization doing the forecasting?

7.2 How exactly can poor forecasts affect an organization's inventory management?

7.3 Many methods of forecasting, and the associated analyses, depend on a knowledge of statistics. General managers do not have this specialized knowledge, so how can they be involved in – or understand the results of – forecasting?

7.4 We have described the most common methods of forecasting. What other methods are available? How can managers choose the best method for their own stocks?

7.5 What are the assumptions of linear regression? Are these generally realistic? How can the method be extended to deal with more complicated situations?

7.6 How can you set the best level of sensitivity for a forecast?

7.7 'Using projective forecasts is like driving a car by looking in the rear view mirror'. What does this mean, and is it true?

References and further reading

Armstrong, J.S. (1995) *Long Range Forecasting*, New York: John Wiley.

De Lurgio, S.A. (1998) *Forecasting Principles and Applications*, New York: Irwin-McGraw Hill.

Diebold, F. (1998) *Elements of Forecasting*, Cincinnati: South Western.

Gaynor, P. and Kirkpatric, R. (1994) *An Introduction to Time Series Modelling and Forecasting for Business and Economics*, London: McGraw-Hill.

Hanke, J.E. and Reitsch, A.G. (1999) *Business Forecasting* (5th edition), Englewood Cliffs, NJ: Prentice Hall.

Pecar, B. (1994) *Business Forecasting for Management*, London: McGraw-Hill.

Thacker, M. (2001a) Why forecast? *Logistics and Transport Focus*, 3(9), 24–29.

Thacker, M. (2001b) Improving forecasting quality, *Logistics and Transport Focus*, 3(10), 60–66.

Wheelwright, S.C., Makridakis, S. and Hyndman, R. (1998) *Forecasting Models for Management* (7th edition), New York: John Wiley.

8

Planning and Stocks

Aims of the chapter

The last two chapters have described the information needed for inventory management. This chapter continues this theme and looks at the information that comes from planned operations. We take the view that stocks are needed to support operations – so to manage the stocks successfully, we need to know about the plans. The chapter shows how plans are designed at different levels, and how these affect stocks. It also lays the foundations for the next chapters on dependent demand methods.

After reading this chapter you should be able to do the following:

- discuss the different levels of planning in an organization and their effect on stocks;
- appreciate the importance of capacity and use a standard approach to capacity planning;
- understand the role of tactical plans, including aggregate plans and master schedules;
- outline methods for designing aggregate plans;
- extend these methods to designing master schedules;
- discuss the role and design of short-term schedules;
- use simulation to describe stock systems.

This chapter emphasizes:

- *plans*, which show how an organization will use its resources;
- *methods of planning*, which describe how the plans are designed at different levels;
- *simulations*, to describe the workings of inventory systems.

Levels of planning

Effects on stock

In the last chapter we saw how to forecast demand using judgemental, projective or causal methods. But there is another approach. Suppose we have a stock of raw materials; if we know what operations are planned to use these materials, then we can get a better idea of how much stock is needed. If, for example, a builder plans to build a certain number of houses over the next three months, it knows how many bricks it will need and when these must be in stock. In other words, it can use planned operations to make decisions about stocks. In this chapter we see how these plans are designed and how they affect stocks.

In Chapter 2 we outlined the hierarchy of decisions that exist in every organization:

- *Strategic decisions* have effects over the long term, involve many resources and are the most risky.
- *Tactical decisions* implement the strategies over the medium term; they look at more detail, involve fewer resources and some risk.
- *Operational decisions* implement the tactics over the short term; they are the most detailed, involve few resources and little risk

The strategic decisions include a broad mission, corporate strategy, business strategy and functional strategies. These are translated into tactical and operational decisions within each function, as shown in Figure 8.1.

A traditional view has senior managers making the strategic decisions that set the organization on its course. These strategic decisions give the objectives, constraints and context for the tactical decisions made by middle managers. These, in turn, give the objectives, constraints and context for operational decisions made by junior managers. To give a specific example, a decision by a manufacturer to open a new factory four or five years in the future is strategic; a decision to modify a production line next year is tactical; a decision about the number of units to make next week is operational. It is difficult to be specific about the timescale or costs of each type of decision as these depend on so many different factors. For electricity generation, a strategic decision might look at building a new power station 20 years into the future and involve expenditure of billions of pounds. For a small restaurant a strategic decision might look one or two years into the future and involve expenditure of a few thousand pounds.

The corporate and business strategies really set the scene for decisions about stocks. There are many alternative strategies, but two dominant ones are 'lean' and 'agile'. A lean strategy does every operation with the least possible resource – people, space, stock, equipment, time, money, etc. An agile strategy gives a high customer service by responding quickly to different or changing circumstances. Under the umbrella of lean and agile strategies, organizations can focus on specific features of their products and operations, such as timing, cost,

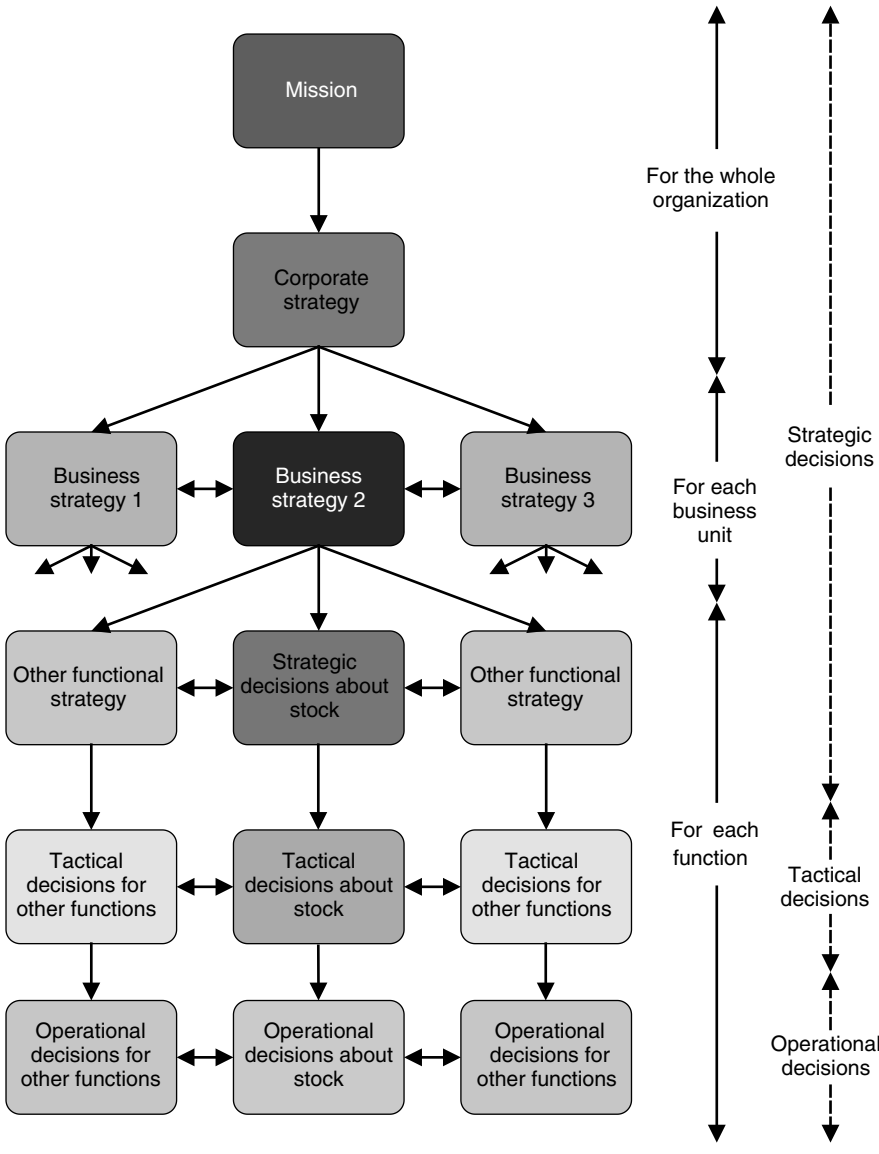


Figure 8.1 Hierarchy of decisions within an organization

quality, product and volume flexibility, diversification, technology, location, and so on.

The important point is that each strategy and focus has a direct effect on stocks. A strategy of low costs will lead to strategic decisions about stock that aim at minimizing long-term costs. But these strategic decisions lead to more detailed tactical plans for the organization – and more detailed tactical decisions about stocks. In turn, the tactical decisions lead to detailed operational decisions

for the organization – and operational decisions about stock. The message is that decisions about stock occur at all levels and are directly linked to other decisions within the organization. If you imagine a strategic issue for stock would be the choice between building warehouses for finished goods or shipping direct to customers from production facilities. This is clearly linked to other decisions within the organization, and it sets the scene for tactical decisions about, say, planned investment in stock during the next few months. These tactical decisions are again linked to other decisions within the organization, and they set the scene for operational decisions about, say, the amount of an item to order this week. To see how this works, we have to follow the planning process down through the organization and look at some of the decisions in more detail.

Capacity planning

A useful place to start looking at planning is with capacity. This sets the maximum amount that can be produced in a given period. The capacity of a delivery system, for example, might be 500 packages a day, while the capacity of a warehouse might be 10,000 pallets. In practice, there is always one resource that is limiting – perhaps the time available on a machine – and this forms a bottleneck that limits overall capacity. This single machine might be working at full capacity, but all the other resources have spare capacity. The kitchens in Hampton Restaurant, for example, can cook 200 meals in an evening, the waiters can serve 250 customers, but the restaurant can only seat 100 customers – so seating is the bottleneck that limits capacity.

With inventory management, our concern is finding the bottleneck in the operations around stocks, and making sure that this does not constrain broader operations. In other words, inventory managers should aim at having enough capacity to meet forecast demand, so that stock is not a bottleneck to the broader operations. The capacity of stocks can be limited by the activities in any of four general areas:

- Capacity of storage, the space available, floor area, investment, number of items, number of bins, length of racking, etc.
- Rate at which materials can be moved into stock, including sorting, checking, break-bulk, unpacking, moving from delivery bays, finding right storage area, handling, etc.
- Rate at which materials can be removed from stock, including locating stock, checking, picking, consolidating orders, moving to departure bays, packing, etc.
- Information handling, including processing of orders, updating stock records, stocktaking, issuing delivery notes, answering queries and any other administration.

The only way of increasing overall capacity is to increase the capacity of the bottleneck. But how much do we want to increase it? The answer to this comes from the demand, as we want the useable capacity to match the demand for products as

closely as possible. Any mismatch between capacity and demand is expensive. If capacity is less than demand, the organization cannot meet all demand and it loses potential customers. If capacity is greater than demand, the organization meets all the demand but it has spare capacity and under-used resources.

We can describe a standard method of overcoming any mismatch between supply and demand. This sees what resources are needed, compares these with the resources available, and then develops plans for overcoming any differences. To be more specific, we must do the following:

1. Examine forecast demand and translate this into a capacity requirement.
2. Find the available capacity of present facilities.
3. Identify mismatches between capacity needed and that available.
4. Suggest alternative plans for overcoming any mismatch.
5. Compare these plans and choose the best.
6. Implement the best.

Worked example

Hazlemere Holdings have forecast the long-term demand for an item at 1,000 units a week. Current facilities have only enough capacity to supply 800 units a week, but Hazlemere can introduce overtime, with fixed costs of £1,000 a week and variable cost of £10 a unit. Alternatively, they can use a sub-contractor who supplies the item for £18 a unit. Present facilities might also be extended, and then each unit could be made for £13. Unfortunately, the extension will take a year to install. What advice would you give to Hazlemere?

Solution

Following the procedure described above, we have the following information:

1. Forecasts show that capacity required is 1,000 units a week.
2. The capacity available is 800 units a week.
3. There is a mismatch between requirements and availability, with a shortage of 200 units a week.
4. The alternatives for overcoming this mismatch are working overtime, using a subcontractor, or extending facilities.
5. We only have some information to compare these alternatives, but can estimate costs of:

working overtime:	$1{,}000 + 200 \times 10$	= £3,000 a week
using a subcontractor:	200×18	= £3,600 a week
extending facilities:	200×13	= £2,600 a week

The best choice is to extend facilities, provided the item can be sold for more than £13 a unit. As this extension will take a year to complete, the company could meet demand until it is finished by working overtime, provided the item can be sold for more than 3,000/200 = £15 a unit.

Unfortunately, taking the steps in this straightforward sequence does not usually work. In most circumstances, there are a huge number of possible plans and we cannot look at all of them in detail. At the same time, it can also be difficult to compare alternatives and pick the best, as there are so many competing objectives and non-quantifiable factors. A more realistic approach replaces the single procedure described with an iterative one. This designs a plan and sees how close it gets to achieving its objectives; if it performs badly, the plan is modified to find improvements. In effect, steps 4 and 5 are repeated until they give a reasonable solution.

We should emphasize two other complications. First, we know that capacity planning is not done in isolation, but is one level of a whole planning process done within an organization. So it has to fit in with higher strategies, and it has to allow feasible lower plans. Second, planning is continuous, and when capacity plans have been finalized for one period, we move on to consider the next period. If we add these two effects, we get the procedure shown in Figure 8.2. This general approach is not restricted to capacity, but can be used for other types of planning.

Capacity planning is primarily a strategic function. When an organization wants to meet demand over the long term, it has to make strategic decisions about the size of buildings, their locations, processes used, levels of technology, and so on. However, when these strategic decisions have been made there is scope for medium- and short-term adjustments to capacity. This might be done by working overtime, using sub-contractors, using an appointment system, rescheduling maintenance, or a host of other options. A manufacturer, for example, might increase capacity by building another factory (which is a strategic decision) and lease additional space while this is being built (which is a tactical decision). While the leased space is being prepared, it might arrange overtime at weekends (which is an operational decision).

Summary

Planning within an organization is done at different levels, which we described as strategic, tactical and operational. Decisions about stock are made at all levels and they are linked to other decisions within an organization at all levels. We saw this effect in capacity planning, where there must be enough capacity for stocks to meet forecast demand, and they should not form a bottleneck that constrains other operations. There is a general procedure for capacity planning, which finds the mismatch between supply and demand, devises plans to overcome this, and implements the best.

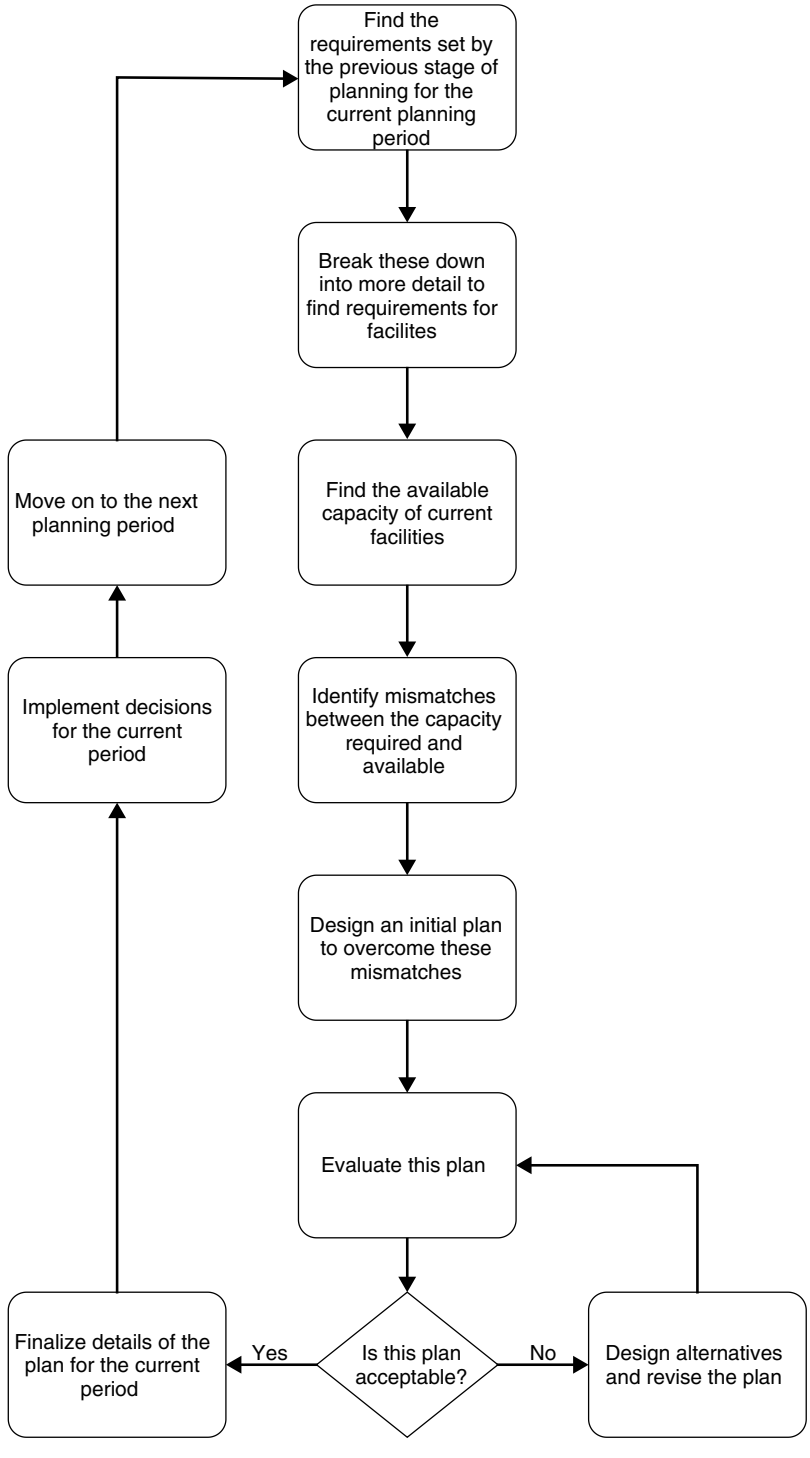

Find the requirements set by the previous stage of planning for the current planning period
Break these down into more detail to find requirements for facilites
Find the available capacity of current facilities
Identify mismatches between the capacity required and available
Design an initial plan to overcome these mismatches
Evaluate this plan
Is this plan acceptable?
Yes
No
Design alternatives and revise the plan
Finalize details of the plan for the current period
Implement decisions for the current period
Move on to the next planning period

Figure 8.2 An iterative planning procedure

Review questions

8.1 Strategic decisions are most important for inventory control. Do you think this is true?
8.2 How are the objectives and constraints set for tactical decisions about stock?
8.3 Why is capacity important for stocks?
8.4 Why is the planning procedure we described iterative?

Aggregate planning

Tactical plans

When strategic plans have been finalized, the next stage is to expand these into tactical plans. There is some disagreement about the terms used to describe these, but we will use the most common:

- *Aggregate plans*, show the overall production planned for families of products, typically by month.
- *Master schedules*, show a detailed timetable of production for individual products, typically by week.

The tactical aggregate plans and master schedules bridge the gap between strategic capacity plans and operational details. So we can combine these into an overall picture and describe the sequence of decisions for, say, a manufacturer planning production.

- The organization's mission gives an overall statement of its aims and beliefs.
- This mission is translated into a corporate strategy that consists of the long-term decisions for the whole organization.
- Each unit within the corporation has a business strategy that includes long-term decisions about products, processes, logistics, facilities, and so on.
- Part of the business strategy includes capacity plans, which match available capacity to long-term demand. Typically, capacity plans set the overall output of each facility by year.
- Capacity plans are expanded into medium-term aggregate plans that, typically, show monthly production of families of products in each facility over the next year or so.
- Aggregate plans are expanded to give master schedules, which show the timetables for production of individual products by week.

- Master schedules are expanded to give short-term schedules that show daily timetables for operations, equipment, people and other resources.

All the decisions that affect stock fit into this general framework of decisions about production.

Aggregate planning

Aggregate plans take the strategic plans, particularly the capacity plans, and add details to give a set of production plans for each family of products, for each of the next few months. Aggregate plans are only concerned with families of products and do not look at individual items. A furniture manufacturer, for example, may produce many different types, styles and sizes of tables and chairs, but the aggregate plan only show the total production of tables and the total production of chairs.

We can summarize the aim of aggregate planning as designing medium-term schedules for making families of products that meet all forecast demand, keep within the constraints of the capacity, use available resources efficiently, keep production relatively stable, and meet any other specific objectives and constraints. To achieve this, the planners have a number of variables they can control. They may, for example, change the product mix, or the number of people employed, the hours worked, the amounts of stock, the amount subcontracted, demand, and so on. Essentially, aggregate planners are looking for answers to a series of questions:

- Should we keep production at a constant level, or change it to meet varying demand?
- Should we adjust the product mix by putting more resources into products with heavier demand?
- Should we use stocks to meet changing demand – producing for stock during periods of low demand and using the accumulated stocks during periods of high demand?
- Should we vary the size of the workforce with demand, hiring or laying-off people?
- Can we change work patterns to meet changing demand by changing shifts or the amounts of overtime and undertime?
- Should we use subcontractors or other outside organizations to cover peak demands?
- Can we have shortages, perhaps with back-orders or late delivery?
- Can we smoothe the demand?

A key question concerns the way that the aggregate plan deals with varying demand. There are basically four alternatives: (1) adjust the supply (by working overtime, etc.), (2) adjust the demand (by changing prices, using incentives, etc.), (3) alter the lead time to customers, and (4) use stocks. Ideally an organization keeps its production constant, as this has benefits of easier planning and control, smoother flow of materials, no problems with changes, employees working regular schedules, no need to 'hire and fire', not using subcontractors, allowing people to become experienced with the operations and able to solve problems, faster throughput, and so on. These benefits mean that organizations generally aim for stable operations, and see major changes in supply as undesirable. At the same time, it can be difficult to adjust the demand or lead time – especially in the short term – as these are largely set by customer expectations and competition. The remaining option, then, is to deal with variable demand by keeping stocks.

Stocks are important for aggregate planning as they separate operations and mean that production in any period need not exactly match demand in the period. In particular, production can remain constant even through periods of varying demand. Then the demand in any period can be met from three internal sources: stocks already held at the beginning of the period, production during the period, or future production with back-orders and late delivery. If we ignore back-orders, we can say that for any period:

$$\begin{array}{l}\text{Stock at end}\\ \text{of this}\\ \text{period}\end{array} = \begin{array}{l}\text{Stock at}\\ \text{end of last}\\ \text{period}\end{array} + \begin{array}{l}\text{Production}\\ \text{during}\\ \text{this period}\end{array} - \begin{array}{l}\text{Demand}\\ \text{met during}\\ \text{this period}\end{array}$$

And if we add the back-orders, but get:

$$\begin{array}{l}\text{Stock at end}\\ \text{of this}\\ \text{period}\end{array} = \begin{array}{l}\text{Stock at}\\ \text{end of last}\\ \text{period}\end{array} + \begin{array}{l}\text{Production}\\ \text{during}\\ \text{this period}\end{array} - \begin{array}{l}\text{Demand}\\ \text{met during}\\ \text{this period}\end{array}$$

$$- \begin{array}{l}\text{Back orders}\\ \text{from earlier}\\ \text{periods}\end{array} + \begin{array}{l}\text{Backorders}\\ \text{met in later}\\ \text{periods}\end{array}$$

Worked example

In April of each year MPG Corp. design next year's production of families of products. Based on last year's figures, the demand for one set of items this year is estimated to be as follows.

Month	1	2	3	4	5	6	7	8	9	10	11	12
Demand	30	50	70	80	90	110	100	80	60	50	50	40

Any shortages are met by back-orders, and there are currently 70 units in stock with back-orders for ten units outstanding. Design an aggregate plan for the year.

Solution

A reasonable initial plan is to keep production constant. The average demand is 67.5 units a month, so we might start by producing 70 units a month to add a bit of safety. The calculations for this are shown in the spreadsheet in Figure 8.3, which has the following entries.

- Column A – the month.
- Column B – monthly demand, with a total of 810 units.
- Column C – opening stock, which is the closing stock for the previous period
- Column D – production, which is initially set at 70 units a month.
- Column E – the back-orders carried over from previous months to be met in this month (which we carry forward from column G).
- Column F – potential shortage, which occurs when the demand to be met (demand plus back-orders from earlier months) is greater than supply (opening stock plus production). If there is a potential shortage, this is met by back orders.
- Column G – back-orders to be met from later production, identified in column F.
- Column H – closing stock, which equals opening stock plus production minus demands met (including adjustments for back-orders).

	A	B	C	D	E	F	G	H
1	**Aggregate planning**							
2								
3								
4	**Month**	**Demand**	**Opening stock**	**Production**	**Back-orders from earlier months**	**Potential shortage**	**Back-orders for later months**	**Closing stock**
5	1	30	70	70	10	0	0	100
6	2	50	100	70	0	0	0	120
7	3	70	120	70	0	0	0	120
8	4	80	120	70	0	0	0	110
9	5	90	110	70	0	0	0	90
10	6	110	90	70	0	0	0	50
11	7	100	50	70	0	0	0	20
12	8	80	20	70	0	0	0	10
13	9	60	10	70	0	0	0	20
14	10	50	20	70	0	0	0	40
15	11	50	40	70	0	0	0	60
16	12	40	60	70	0	0	0	90
17								
18	**Totals**	**810**		**840**				**830**

Figure 8.3 Initial aggregate plan for worked example

The next stage would be to assign costs for stocks, shortage, etc. and see how well this plan meets the company's objectives. An obvious problem is that stocks are high in the first few months. We could get around this by reducing production to, say, 30 units for the first three months to use existing stock, and then raise it to 80 units for the remaining months. The results for this are shown in Figure 8.4. As you can see, this reduces stocks, but there are a lot of back-orders during the middle of the year. This may not be acceptable, so we need to improve our plan and keep adjusting it until we get a reasonable solution.

	A	B	C	D	E	F	G	H
1	**Aggregate planning**							
2								
3								
4	**Month**	**Demand**	**Opening stock**	**Production**	**Back-orders from earlier months**	**Potential shortage**	**Back-orders for later months**	Closing stock
5	1	30	70	30	10	0	0	60
6	2	50	60	30	0	0	0	40
7	3	70	40	30	0	0	0	0
8	4	80	0	80	0	0	0	0
9	5	90	0	80	0	10	10	0
10	6	110	0	80	10	40	40	0
11	7	100	0	80	40	60	60	0
12	8	80	0	80	60	60	60	0
13	9	60	0	80	60	40	40	0
14	10	50	0	80	40	10	10	0
15	11	50	0	80	10	0	0	20
16	12	40	20	80	0	0	0	60
17								
18	**Totals**	**810**		**810**				**180**

Figure 8.4 Modified aggregate plan for worked example

Methods of designing aggregate plans

In the last worked example we chose an initial plan and used intuition to look for improvements. In practice, there are many ways of designing plans, ranging from simple intuition through to sophisticated mathematical models. The most appropriate depends on a number of factors, ranging from skills available within the organization to the cost of production. If you run a large oil company, production costs are very high and it is worth putting a lot of effort into a sophisticated mathematical programming model that gets very good results. On the other hand, a small business is unlikely to have the resources for this, and could not justify the extra effort to get marginally better results. The following list gives the most important methods of designing plans:

1. *Negotiations*. Planning is so complicated, with many subjective factors and people affected, that the best approach is often to negotiate a solution. This may not give the best technical answer, but it has the support of everyone concerned.

2. *Adjust previous plans*. The best operations are relatively stable, so a useful approach to planning has an experienced planner reviewing previous plans and updating them to allow for any changing circumstances. This has the benefit of being relatively easy, causing little disruption, using a well-understood procedure and results that are trusted by the organization. Unfortunately, the results can also be of variable and uncertain quality, the plans may take a long time to design, and they rely solely on the skills of a planner.

3. *Other intuitive methods*. These include a range of methods that use the skills, knowledge and experience of planners, who will typically use a series of heuristic rules that have been successful in the past.

4. *Graphical methods*. Planners often find it easier to work with graphs or diagrams. The most popular format uses a graph of cumulative demand over some time period, and the corresponding line of cumulative supply. The aim is to get the cumulative supply line nearly straight – giving constant production – and as close as possible to the cumulative demand line. The difference between the two lines shows a mismatch:

   - if the cumulative demand line is below the cumulative supply line, production is too high and the excess is accumulated as stock;
   - if the cumulative demand line is above the cumulative supply line, production has been too low and demand is not being met.

   Graphical approaches have the advantages that they are easy to use and understand. But they are really only one step better than an intuitive method. They do not guarantee optimal solutions, sometimes give very poor results, may take a long time, and still rely on the skills of a planner.

5. *Other spreadsheet calculations*. In the last method, we concentrated on the graphs, and probably used a spreadsheet to do related calculations. Concentrating on the graphs might show the overall patterns, but it might lose some of the details. An alternative is to concentrate on the spreadsheet calculations and look at the patterns in the numbers.

6. *Simulation*. Simulation is one of the most flexible approaches to solving problems. It gives a dynamic view by imitating real operations over a typical period. We will return to this approach later in the chapter.

7. *Expert systems*. These specialized programs try to make computers duplicate the thinking of a skilled scheduler. The basic skills, expertise, decisions and rules used by experts are collected in a knowledge base. A user of the system passes a specific problem to an inference engine, which controls the rules used to get a solution mechanism.

8. *Mathematical models*. The methods we have listed so far rely, at least to some extent, on the skills of a planner. More formal mathematical approaches give optimal – or near optimal – solutions without any human intervention. In practice, aggregate plans include so many subjective and non-quantifiable factors, that optimal solutions in the mathematical sense may not give the best answers for the organization.

   The most common mathematical approach uses linear programming, which we can illustrate in the following problem. These methods are fairly complicated, and are generally limited to smaller problems. If, however, you work in a chemical plant that produces millions of litres a month, a small change in the aggregate plans can make a significant difference to costs, and it is certainly worth looking at mathematical approaches.

Worked example

The main production costs incurred by a company are:

- supplying a unit of product;
- holding stocks;
- every unit of unmet demand;
- amount of overtime used;
- amount of undertime (i.e. normal working time that is not used);
- increase in production rate;
- decrease in production rate.

If the immediate objective is to minimize total production costs, how would you formulate this as a linear programme?

Solution

We can start by defining the costs as:

C_V = variable cost of supplying a unit

C_H = cost of holding a unit of stock for a unit time

C_S = shortage cost per unit of unmet demand

C_O = additional cost per unit made with overtime

C_U = cost per unit of undertime

C_I = cost of increasing the production rate

C_R = cost of reducing the production rate.

There are two other constants:

$$D_t = \text{demand in period t}$$

$$N_t = \text{normal capacity in period t.}$$

Now we can define the variables:

$$P_t = \text{production in period t}$$

$$H_t = \text{stock held at the end of period t}$$

$$S_t = \text{shortage, or unmet demand in period t}$$

$$O_t = \text{units produced by overtime in period t}$$

$$U_t = \text{units of undertime in period t}$$

$$I_t = \text{increase in production rate during period t}$$

$$R_t = \text{reduction in production rate during period t.}$$

With these values we can define the objective function as minimizing the total cost:

$$\text{Minimize } \{C_V P_t + C_H H_t + C_S S_t + C_O O_t + C_U U_t + C_I I_t + C_R R_t\}$$

There are a number of constraints that hold for every period.

Stock holdings must balance the supply, demand and shortages:

$$H_t = H_{t-1} + P_t - D_t + S_t$$

Total production equals normal production plus overtime minus undertime:

$$P_t = N_t + O_t - U_t$$

Changes in production rates are consistent with production in each period:

$$P_t - P_{t-1} = I_t - R_t$$

Now we could substitute the known values and use a linear programming package to find the optimal solution for production. As you can see, though, this is a very simple version of a real problem, but a 12-month period still needs accurate values for 7 costs, 24 constants, 84 variables and 36 constraints. Problems of any real size and complexity soon become unwieldy. So linear programming has the disadvantage of being complicated and needing a lot of effort, skills and experience. It is also difficult to understand, is time-consuming, is expensive, needs a lot of reliable data, and the model still may not be a good description of the real situation. On the other hand, it has the advantage of guaranteeing an 'optimal' solution.

Worked example

Jane Allison has recently forecast monthly demand for a family of products, as shown below. At the end of each month she assigns a notional storage cost of £20 for every unit held in stock. Any shortages are satisfied by back-orders, but each unit of shortage is assigned a notional cost of £200 for lost profit, goodwill and future sales. Each time the production rate changes it costs £10,000. The capacity for the products is 600 units a month. How can Jane design an aggregate plan for the products?

Month	1	2	3	4	5	6	7	8	9
Demand	560	640	520	320	240	200	120	200	260

Solution

Jane can use a graphical method here and will probably start with a stable production to avoid the expensive changes. Setting production at the average demand of 340 units a month gives the cumulative demand and supply shown in Figure 8.5. The cumulative demand line is always above the cumulative supply. There are no storage costs, but the continuous shortages give high overall costs.

Changing the production rate is expensive, but it might be worthwhile to reduce the shortages in the initial plan. As demand is heavy in the first three months, Jane might try increasing supply by running the process at its maximum output of 600 units a month. The total demand to be met from production in the remaining six months is $(3{,}060 - 3 \times 600) = 1{,}260$, averaging 210 a month. Figure 8.6 shows that there is clearly a better match between the graphs of cumulative supply and demand and the costs are lower. Jane could continue adjusting this plan until she got a result that achieved her objectives.

Summary

Strategic plans are followed by more detailed tactical and operational plans. We can describe tactical plans as aggregate plans (typically showing monthly production of families of products) and master schedules (typically showing weekly production of each product). The main aims of aggregate planning are to meet the requirements of higher plans while using resources efficiently. Stocks play an important role in the design of aggregate plans, as this means that production does not have exactly to match demand. There are many ways of designing aggregate plans, ranging from intuition to mathematical models.

	A	B	C	D	E	F	G	H
1	**Aggregate planning**							
2								
3	**Month**	**Aggregate demand**	**Cumulative demand**	**Production**	**Cumulative production**	**Shortage in month**	**Stock at end of month**	
4	**1**	560	560	340	340	220	0	
5	**2**	640	1200	340	680	520	0	
6	**3**	520	1720	340	1020	700	0	
7	**4**	320	2040	340	1360	680	0	
8	**5**	240	2280	340	1700	580	0	
9	**6**	200	2480	340	2040	440	0	
10	**7**	120	2600	340	2380	220	0	
11	**8**	200	2800	340	2720	80	0	
12	**9**	260	3060	340	3060	0	0	
13								
14	**Totals**					**3440**	**0**	
15								
16	**Costs**		**Production changes**			**1**	**£10,000.00**	
17			**Shortages**			**3,440**	**£688,000.00**	
18			**Storage**			**0**	**£0.00**	
19			**Total**				**£698,000.00**	
20			**Monthly**				**£77,555.56**	
21								
22								
23								
24								
25								
26								
27								
28								
29								
30								
31								
32								
33								
34								
35								
36								
37								
38								
39								
40								
41								
42								

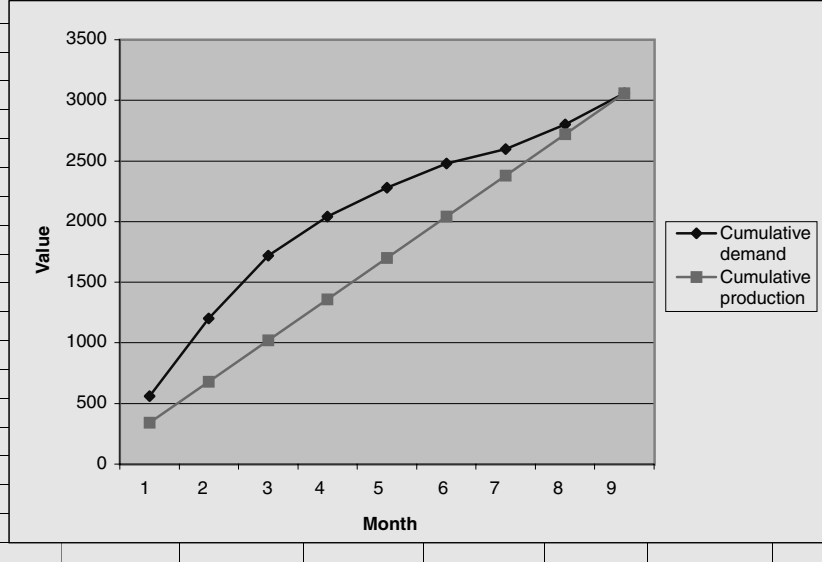


Figure 8.5 Initial aggregate plan for worked example

	A	B	C	D	E	F	G	H
1	**Aggregate planning**							
2								
3	**Month**	**Aggregate demand**	**Cumulative demand**	**Production**	**Cumulative production**	**Shortage in month**	**Stock at end of month**	
4	**1**	560	560	600	600	0	40	
5	**2**	640	1200	600	1200	0	0	
6	**3**	520	1720	600	1800	0	80	
7	**4**	320	2040	210	2010	30	0	
8	**5**	240	2280	210	2220	60	0	
9	**6**	200	2480	210	2430	50	0	
10	**7**	120	2600	210	2640	0	40	
11	**8**	200	2800	210	2850	0	50	
12	**9**	260	3060	210	3060	0	0	
13								
14	**Totals**					**140**	**210**	
15								
16	**Costs**		**Production changes**			**1**	**£10,000.00**	
17			**Shortages**			**140**	**£28,000.00**	
18			**Storage**			**210**	**£4,200.00**	
19			**Total**				**£42,200.00**	
20			**Monthly**				**£4,688.89**	
21								
22								
23								
24								
25								
26								
27								
28								
29								
30								
31								
32								
33								
34								
35								
36								
37								
38								
39								
40								
41								
42								

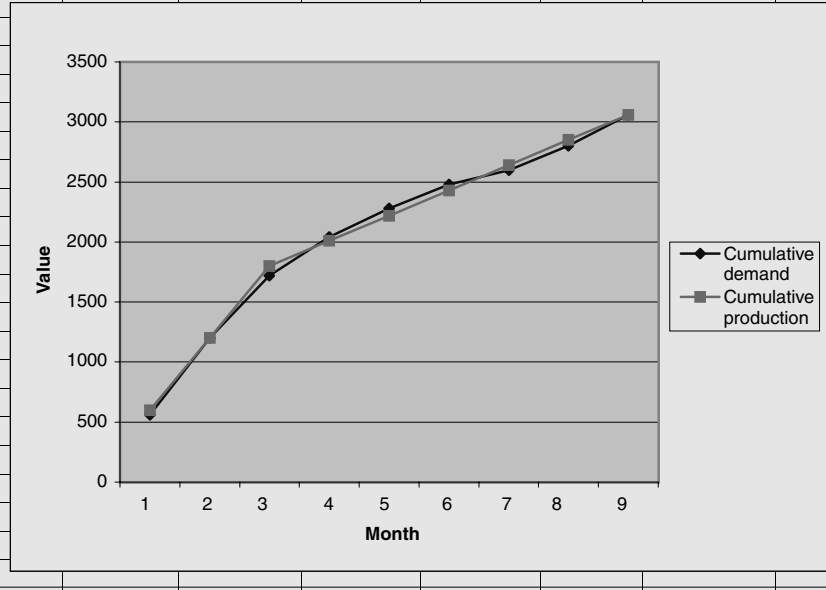


Figure 8.6 Modified aggregate plan for worked example

Review questions

8.5 How often would a typical organization update its aggregate plans for a particular product?
8.6 What is the main output from aggregate planning?
8.7 Why are stocks important for aggregate planning?
8.8 What are the benefits of using spreadsheets for aggregate planning?
8.9 How can you recognize a good aggregate plan from a graph?
8.10 When would you use mathematical programming for aggregate planning?

Master schedules

Expanding the aggregate plans

We can expand the aggregate plan to give more details, which are shown in a *master schedule*. The master schedule 'disaggregates' the aggregate plan and shows the number of individual products to be made, typically, each week. This detailed timetable of planned production shows the due dates or completion times of individual products. An aggregate plan may, for example, show 1,000 motors being made next month, while the master schedule shows 100 standard outboard motors and 150 5 kW electric motors in week 1, 50 standard outboard motors and 200 3 kW electric motors in week 2, and so on. This is the first time that due dates are associated with individual products.

The master schedule is derived from the aggregate plan, so the total production given in the master schedule must equal the production specified in the aggregate plan. In practice, master schedules are designed close to the start of production, so they need to be more flexible. There may be, for example, some adjustments to allow for orders already received, back-orders, revised forecasts, new constraints on operations, and so on. If everything has been done properly, these effects should be small, and master schedules should be derived directly from the aggregate plans. Then, the overall objective of the master schedule is to devise a detailed timetable for individual products that allows the aggregate plan to be achieved as efficiently as possible.

Figure 8.7 shows part of a simplified master schedule. Rows 3 to 5 show the quarterly and monthly production targets from the aggregate plan. Rows 8, 9 and 10 show the master schedule for three products to meet these targets. Rows 11 and 12 confirm that these targets match the aggregate plan.

Rows 14 to 18 review the current orders for one item. Row 14 shows the planned monthly production from row 8. Row 15 shows the orders already received, and row 16 shows an intermediate stage where production is 'committed', perhaps for orders that are expected but still being negotiated. Rows 17 and 18 show the units allocated to this known demand. The planned production is higher than current known demand, but as time passes the two should come into line, with short-term variations met by changing stock levels. The complete master schedule would give more details of the plans for this item, and then move on to all other items.

	A	B	C	D	E	F	G	H	I	J	K
1	**Master schedule**										
2											
3	**Quarter**	3	**Plan**	5020							
4	**Month**	**April**				**May**				**June**	
5	**Production**	1200				2080				1740	
6											
7	**Week**	**1**	**2**	**3**	**4**	**1**	**2**	**3**	**4**	**1**	**2**
8	**Part 1975389**	200	120	60	0	0	0	150	400	150	150
9	**Part 925635**	100	80	40	0	0	100	300	300	300	300
10	**Part 2354672**	0	100	200	300	400	400	30	0	0	0
11	**Totals**	300	300	300	300	400	500	480	700	450	450
12	**Planned**	1200				2080				1740	
13											
14	**Part 1975389**	380				550				350	
15	**Orders**	120	80	30	20	6	1	0	0	0	0
16	**Committed**	12	14	20	14	8	6	6	1	1	0
17	**Totals**	132	94	50	34	14	7	6	1	1	0
18	**Allocated**	310				28				2	

Figure 8.7 Example of a master schedule

Designing master schedules

In principle, designing the master schedule is similar to designing an aggregate plan, so methods again range from discussion through to mathematical models. As we are looking at lower levels and more details, plans become increasingly complicated and messy. They can involve a lot of subjective views, so master schedules are more likely to be designed by skilled schedulers using some intuitive approach.

Worked example

FNT Outdoor Electronics plans to produce 2,400 units in June and 2,000 units in July in one of its factories. Now they need to design a master schedule for their two main products, a high powered amplifier and a stadium display unit. They already have stocks of 140 units of the amplifier and 100 units of the display unit, and the factory has a capacity of 600 units a week. Sales of the amplifier are normally twice as high as sales of the display units and FNT already have orders for:

Week	1	2	3	4	5	6
Amplifier	510	360	340	220	160	0
Display unit	270	210	150	80	40	0

Solution

The master schedule must allow the production of 2,400 units in the first month, and the maximum capacity is 600 units a week. Two-thirds of demand is usually

for amplifiers, so an obvious plan has production of 600 units a week, with 400 amplifiers and 200 display units. During the second month the aggregate plans calls for 2,000 units, so an obvious plan is to make 500 units a week, 333 amplifiers and 167 display units.

This plan has to be adjusted in the short term, because there are already some orders and stocks, so FNT could start by meeting all existing orders, and then assign two-thirds of spare production to making amplifiers and one-third to making display units. There are already orders of 780 units in week 1, so the first priority is to meet these. Orders for 140 amplifiers and 100 display units can be met from stock, with the remaining 540 units met from production. This leaves 60 units of free production which is divided into 40 units amplifiers and 20 display units.

In week 2 the demand for 360 amplifiers and 210 display units is met partly from stock (60 units) and partly from production (510 units). This leaves 90 units of free production which is divided into 60 amplifiers and 30 display units. Repeating this procedure for every period gives the results shown in Figure 8.8. The build-up of stock in later weeks shows that this production has not yet been allocated to customers and shows the stock levels if no more orders are received.

This solution is, of course, only one of many possible ones. It has the advantages of meeting the aggregate plan and keeping production stable, but FNT could now start looking for improvements. They could, of course, also use a more formal approach than this intuitive one.

	A	B	C	D	E	F	G	H	I
1	**Master schedule**								
2									
3	**Month**	**June**				**July**			
4	**Week**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
5									
6	***Amplifier***								
7	**Actual orders**	510	360	340	220	160	0	0	0
8	**Opening stock**	140	40	60	133	333	533	867	1200
9	**Requirement**	370	320	280	87	–173	–533	–867	–1200
10	**Production**	410	380	413	420	360	333	333	333
11									
12	***Display unit***								
13	**Actual orders**	270	210	150	80	40	0	0	0
14	**Opening stock**	100	20	30	67	167	267	433	600
15	**Requirement**	170	190	120	13	–127	–267	–433	–600
16	**Production**	190	220	187	180	140	167	167	167
17									
18	***Total***								
19	***Orders***	780	570	490	300	200	0	0	0
20	***Production***	600	600	600	600	500	500	500	500
21	**Aggregate plan**	2400				2000			

Figure 8.8 Master schedule for worked example

Summary

The master schedule disaggregates the aggregate plan to give a detailed timetable for making each product, typically each a week. Designing a master schedule is similar in principle to designing an aggregate plan, and similar methods can be used. The subjective decisions and timing mean that simpler methods are more likely to be used for master schedules.

Review questions

8.11 What is the main purpose of a master schedule?

8.12 What are the main constraints on a master schedule?

Operational schedules

The iterative adjustment of the master schedules can continue up to some time – say three weeks – before production actually starts. After this the schedule must be finalized to allow the more detailed planning of equipment, materials, people and other resources. There are two ways that we can approach this more detailed, operational scheduling. First, we can continue with the method that we have used for higher plans, comparing requirements and availability of resources, designing plans to overcome any mismatches, and so on. This gives *short-term scheduling*, which we will discuss here. The other method takes a fundamentally different view of dependent demand, and we will describe this in following chapters.

Short-term scheduling

The master schedule shows the production planned in, say, the next week, and this sets the requirements and constraints for the most detailed, operational level of planning. If we want to make 100 units next week, we can start by using four machines and two operators on Monday morning. All resources used in an organization eventually need this kind of detailed schedule.

Essentially, the methods available for short-term scheduling are the same as those for aggregate panning and master scheduling, and they range from discussion through to mathematical modelling. With short-term scheduling, though, there is little time available, and we have to be more flexible to deal with changing circumstances and requirements. The consequences of using schedules that are not 'optimal' are generally less severe than at higher levels. As a result, the usual methods of short-term scheduling are much more likely to be intuitive and using heuristics.

The most common approach to short-term scheduling uses simple *scheduling rules*. These are guidelines or rules of thumb that seem to give reasonable results. We can illustrate the approach for processing batches of different items on a piece of equipment. Then the object of scheduling is to design a timetable that allows the batches to move as efficiently as possible through the equipment (as

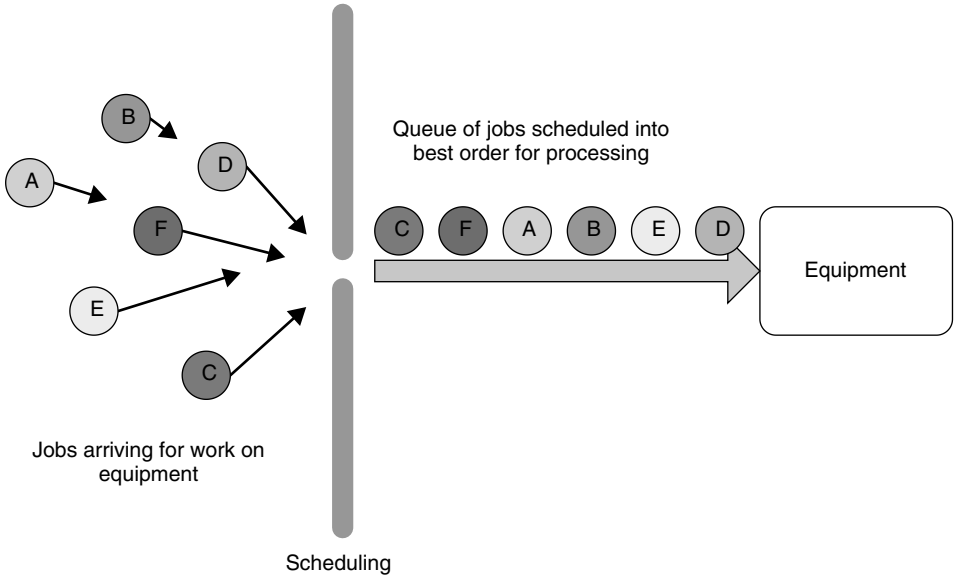


Figure 8.9 Example of short-term scheduling

shown in Figure 8.9). In this context, we can think of efficiency in several different ways, perhaps minimizing the waiting time of batches, minimizing the total processing time, keeping stocks low, reducing the maximum lateness, achieving high utilization of equipment, or some other specific measure. Then we can use different rules to achieve each of these objectives. One thing that we cannot do is change the total time for processing all batches, as this remains constant regardless of the order in which we schedule the jobs (assuming that the set-up time for each batch is fixed regardless of the batch that preceded it). But different schedules do affect other measures of performance. Some common rules for designing these include the following:

1. *First-come-first-served* is the most obvious scheduling rule and simply schedules batches in the order they arrive. It assumes no priority, no urgency, or any other measure of relative importance. The benefits of this rule are simplicity, ease of use and clear equity. The drawbacks are that urgent or important jobs may be delayed while less urgent or important ones are being processed.
2. *Most urgent first* assigns an importance, or urgency, to each batch and they are processed in order of decreasing urgency. A hospital emergency department, for example, will schedule patients in this order. The benefit of this rule is that more important batches are given higher priority, but jobs with low priority might be stuck at the end of a queue for a very long time.
3. *Shortest job first* minimizes the average time spent in the system, where this time is defined as processing time plus waiting time. A batch that can be finished quickly is processed and moved on through the system, while longer jobs are

left until later. The disadvantage of this rule is that long jobs can spend a long time in the system.

4. *Earliest due date first* takes batches in the order that they are needed for customers. This has the benefit of minimizing the maximum lateness of batches, but again some batches may have to wait a long time.

5. *Critical ratio* The critical ratio is the time remaining until a batch is needed divided by the time needed to complete the batch. If this ratio is low, the time to complete the batch is short compared with the time available and the batch becomes urgent: if the ratio is high, there is plenty of time remaining and the batch is less urgent. This ratio changes as batches move through operations and priorities depend on the progress of other jobs.

This list suggests some simple scheduling rules, but there are many more that cover a range of different objectives and circumstances. We could, for example, schedule batches in the order of least work remaining, or fewest operations remaining; we might look at the next process and consider combined times for two or three of these; we could look at the slack (which is the time remaining until the job is due minus the time remaining for processing). Each of these – and the many alternatives – could be useful in particular circumstances, and the choice of the best must be a matter of judgement.

Worked example

Six batches of items are to be scheduled on a piece of equipment. Each batch replaces parts used in production, and fully occupies the equipment for the times given below. If the average demand for products and current stocks are as follows, what schedule would you suggest?

Batch	1	2	3	4	5	6
Duration (days)	9	5	2	4	1	4
Demand	3	5	13	1	5	25
Current stock	84	65	260	7	25	225

Solution

Each batch has a different priority that depends on the time when stock of the item will run out. We can, then, schedule batches in the order most-urgent-first, where urgency is measured by the number of days of stock remaining stock or current stock divided by mean demand. This gives the result shown in Figure 8.10. As you can see, all batches are finished before the items are due to run out, except for batch 2, where stocks run out on day 13 but the batch is finished on day 14. An obvious way around this problem is to have part of batch 2 delivered a day early. Figure 8.10 also shows the results of two other scheduling rules for comparison.

	A	B	C	D	E	F	G
1	**Short-term schedules**						
2							
3	**Data**						
4	Batch	1	2	3	4	5	6
5	Duration	9	5	2	4	1	4
6	Stock remaining	84	65	260	7	25	225
7	Demand	3	5	13	1	5	25
8	Days stock remaining	28	13	20	7	5	9
9							
10	**Most urgent first**						
11	Batch	5	4	6	2	3	1
12	Start	0	1	5	9	14	16
13	Duration	1	4	4	5	2	9
14	Finish	1	5	9	14	16	25
15							
16	**First-come-first-served**						
17	Batch	1	2	3	4	5	6
18	Start	0	9	14	16	20	21
19	Duration	9	5	2	4	1	4
20	Finish	9	14	16	20	21	25
21							
22	**Shortest first**						
23	Batch	5	3	4	6	2	1
24	Start	0	1	3	7	11	16
25	Duration	1	2	4	4	5	9
26	Finish	1	3	7	11	16	25

Figure 8.10 Different scheduling rules for worked example

Control of schedules

Short-term schedules show what each job, piece of equipment, person and every other resource should be doing at any time. But there is a difference between designing timetables and having them actually happen. There are often unforeseen factors that prevent the plans actually happening, such as equipment developing faults, people being ill, suppliers not sending the right materials, late deliveries, and so on. To keep a check on such problems, and find ways of overcoming them, organizations use a *control system*.

Control systems have two main parts. The first part checks the progress of jobs and gives information back to managers. This part is always working, monitoring operations and making sure that things are going smoothly. It checks details of progress, reports the actual times of operations, records measure of performance (such as efficiency, productivity and utilization), and reports any problems. The second part starts working when circumstances change, or something goes wrong. Then plans have to be revised, and the control system either warns about the

changes needed, or it may actually adjust the schedules. To be more specific, the purpose of a control system is to do the following:

- monitor operations and report on their performance;
- make sure jobs are scheduled according to plans;
- warn of problems with resources, delivery dates, etc.;
- check progress as jobs move through the process;
- make small adjustments as necessary to plans;
- allow rescheduling if there is a major disruption to plans;
- give information on current activities.

Summary

Short-term scheduling designs detailed timetables for every piece of equipment and resource. Because of the difficulty, timing, flexibility and subjective decisions, this is usually done by intuitive methods and heuristics. Scheduling rules are commonly used, and there are many alternatives for different circumstances. A control mechanism is needed to compare actual schedules with plans, and make necessary adjustments.

Review questions

8.13 What is the aim of short-term scheduling?

8.14 Why would you use scheduling rules when more sophisticated methods are available?

8.15 What is the purpose of a control system for schedules?

Simulation of stocks

Managers continually monitor their operations and look for improvements. But when they have an idea for a change, how do they know that the result will actually be an improvement and there will not be some unforeseen consequences? They clearly need some way of studying the effects of their decisions. If, for example, they decide to lower the stocks of an item, they want to know the effects, not only on stocks but on associated operations. They could, of course, experiment with real operations, reducing the stock level and seeing what happens. If things go wrong, though, this could be disruptive and expensive. As an alternative they could use mathematical models. Unfortunately, stock control would need very complicated models that include a lot of simplifications, approximations and assumptions. Even then, analytical models are simplifications of reality, and they may deal badly with some important factors.

Another alternative is to use *simulation*. This is based on a dynamic representation that duplicates the continuous operation of a system over some time. An ordinary model for inventory control looks at the operations at a specific time. It

describes a situation at that point, collects data relevant to that time and draws conclusions. Simulation on the other hand follows the operations and sees exactly what happens over some extended time. A simple analogy has an ordinary model giving a snapshot of the operations, while simulation takes a movie.

Building a simulation model

Simulation is based on an appropriate model. The first step in building a simulation model is to describe the operations in detail. There are several ways of doing this, with the most common using a flow diagram show the sequence of activities. Suppose, for example, we want to find the effects of using an economic order quantity, when we know the monthly demand for an item. We see how the stock operations work, and we might get the procedure shown in Figure 8.11.

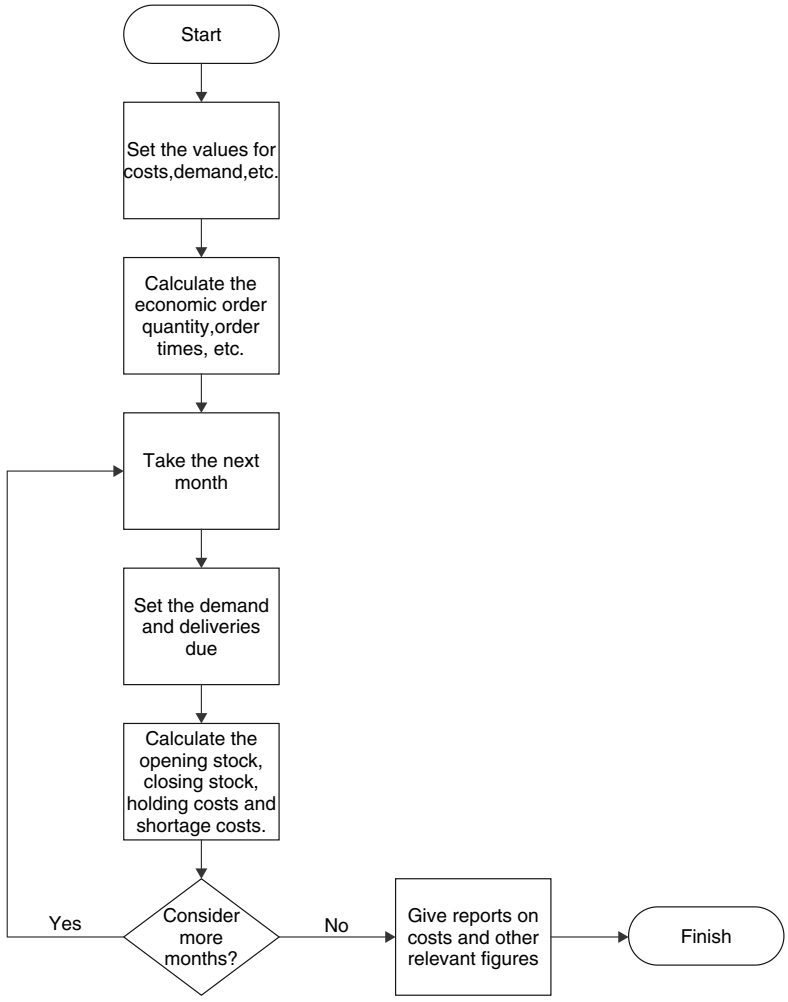

Start
Set the values for costs,demand,etc.
Calculate the economic order quantity,order times, etc.
Take the next month
Set the demand and deliveries due
Calculate the opening stock, closing stock, holding costs and shortage costs.
Consider more months?
Yes
No
Give reports on costs and other relevant figures
Finish

Figure 8.11 Steps in a simple inventory simulation

This starts by setting the known data, including costs and demand pattern, and uses this to calculate the economic order quantity, timing of deliveries, and so on. Then we take the first month, check the demand and deliveries due, find the opening stock, set the opening stock (which is last month's closing stock), find the closing stock (which is opening stock plus deliveries minus demand) and calculate all the costs. Then we go to the second month and repeat the analysis. We can continue this for as many repetitions as needed to get a longer-term view of the operations. Then we can summarize all the figures and check the effect of our economic order quantity. All the calculations for this are done with appropriate software, and Figure 8.12 shows how they might appear in a spreadsheet. In this example the economic order quantity is calculated as 200, corresponding to one delivery a month. The simulation follows the operations for 12 months, as shown in rows 21 to 32. As you can see, the costs are high as there is a build-up of stock in

	A	B	C	D	E	F	G
1	**Simulation of stocks**						
2							
3	**Notes:**	Effects of variable demand on costs					
4							
5	**Inputs**	Annual demand		2400			
6		Unit cost		£ 10.00			
7		Holding cost amount		£ 6.00			
8		Holding cost percentage		0%			
9		Reorder cost		£ 50.00			
10		Shortage cost		£ 100.00			
11							
12	**Results**	Annual holding cost		£ 6.00			
13		Order size		200.00			
14		Time between orders		1.00			
15		Number of orders a year		12.00			
16		Fixed unit costs		£ 24,000.00			
17		Variable cost		£ 1,200.00			
18		Total annual cost		£ 25,200.00			
19							
20	**Month**	**Opening stock**	**Demand**	**Delivery**	**Closing stock**	**Holding cost**	**Shortage cost**
21	1	0	150	200	50	£ 25.00	£ -
22	2	50	120	200	130	£ 65.00	£ -
23	3	130	150	200	180	£ 90.00	£ -
24	4	180	150	200	230	£ 115.00	£ -
25	5	230	300	200	130	£ 65.00	£ -
26	6	130	250	200	80	£ 40.00	£ -
27	7	80	300	200	–20	£ -	£ 166.67
28	8	–20	200	200	–20	£ -	£ 166.67
29	9	–20	250	200	–70	£ -	£ 583.33
30	10	–70	220	200	–90	£ -	£ 750.00
31	11	–90	200	200	–90	£ -	£ 750.00
32	12	–90	110	200	0	£ -	£ -
33							
34					Summary costs		
35						Holding cost	£ 400.00
36						Reorder cost	£ 600.00
37						Shortage cost	£ 2,416.67
38						Unit cost	£ 24,000.00
39						Total cost	£ 27,416.67

Figure 8.12 Simulating the effects of variable demand with the economic order quantity

the first half of the year, but consistent shortages in the second half. This suggests that we might look for some other pattern of ordering, and we could check its effect in our simulation model.

Now we have the basic steps in a simulation, we can start adding some more details to give a better picture of the real system. We could, for example, vary the pattern of demand rather than consider one fixed pattern, and we could add a variable lead time and amount delivered (perhaps allowing for quality checks). The way of introducing this kind of variation is to use random numbers. Suppose that you do not know exactly the demand in a month, but know that it is between, say, 30 and 50 units. Spreadsheets can automatically generate random values and in Microsoft Excel the function RANDBETWEEN(30,50) will generate a random value between 30 and 50. Then for each subsequent month it will generate a new value in this range. Figure 8.13 shows a modification to the previous spreadsheet which allows variation in the demand, amount delivered and lead time and also

	A	B	C	D	E	F	G
1	**Simulation of stocks**						
2							
3	**Notes:**	Effects of variable demand on costs					
4							
5	**Inputs**	Annual demand		2400			
6		Unit cost		£ 10.00			
7		Holding cost amount		£ 6.00			
8		Holding cost percentage		0%			
9		Reorder cost		£ 50.00			
10		Shortage cost		£ 100.00			
11							
12	**Results**	Annual holding cost		£ 6.00			
13		Order size		200.00			
14		Time between orders		1.00			
15		Number of orders a year		12.00			
16		Fixed unit costs		£ 24,000.00			
17		Variable cost		£ 1,200.00			
18		Total annual cost		£ 25,200.00			
19							
20	**Month**	**Opening stock**	**Demand**	**Delivery**	**Closing stock**	**Holding cost**	**Shortage cost**
21	1	0	193	196	3	£ 1.50	£ -
22	2	3	158	200	45	£ 25.50	£ -
23	3	44	203	198	39	£ 19.50	£ -
24	4	39	232	208	15	£ 7.50	£ -
25	5	15	193	204	26	£ 13.00	£ -
26	6	26	227	198	–3	£ -	£ 25.00
27	7	–3	237	189	–51	£ -	£ 425.00
28	8	–52	165	195	–22	£ -	£ 183.33
29	9	–23	216	182	–57	£ -	£ 475.00
30	10	–57	178	195	–40	£ -	£ 333.33
31	11	–40	217	196	–61	£ -	£ 508.00
32	12	–61	183	184	–60	£ -	£ 500.00
33							
34					**Summary costs**		
35						Holding cost	£ 64.00
36						Reorder cost	£ 600.00
37						Shortage cost	£ 2,450.00
38						Unit cost	£ 24,000.00
39						**Total cost**	£ 27,114.00

Figure 8.13 Adding some random variation to the simulation

allows for the odd unit to go missing from stock. This gives the results from one simulation run, but the random variations mean that each run will be different. To find the real patterns we would have to repeat the simulation many times, perhaps hundreds or even thousands. However, once we have built the model this is quite easy, particularly using one of the many standard simulation packages.

We have added some features to our original simulation model, but we could continue doing this, expanding the model in many ways to suit individual circumstances. We can illustrate the start of one expansion in Figure 8.14. The operations

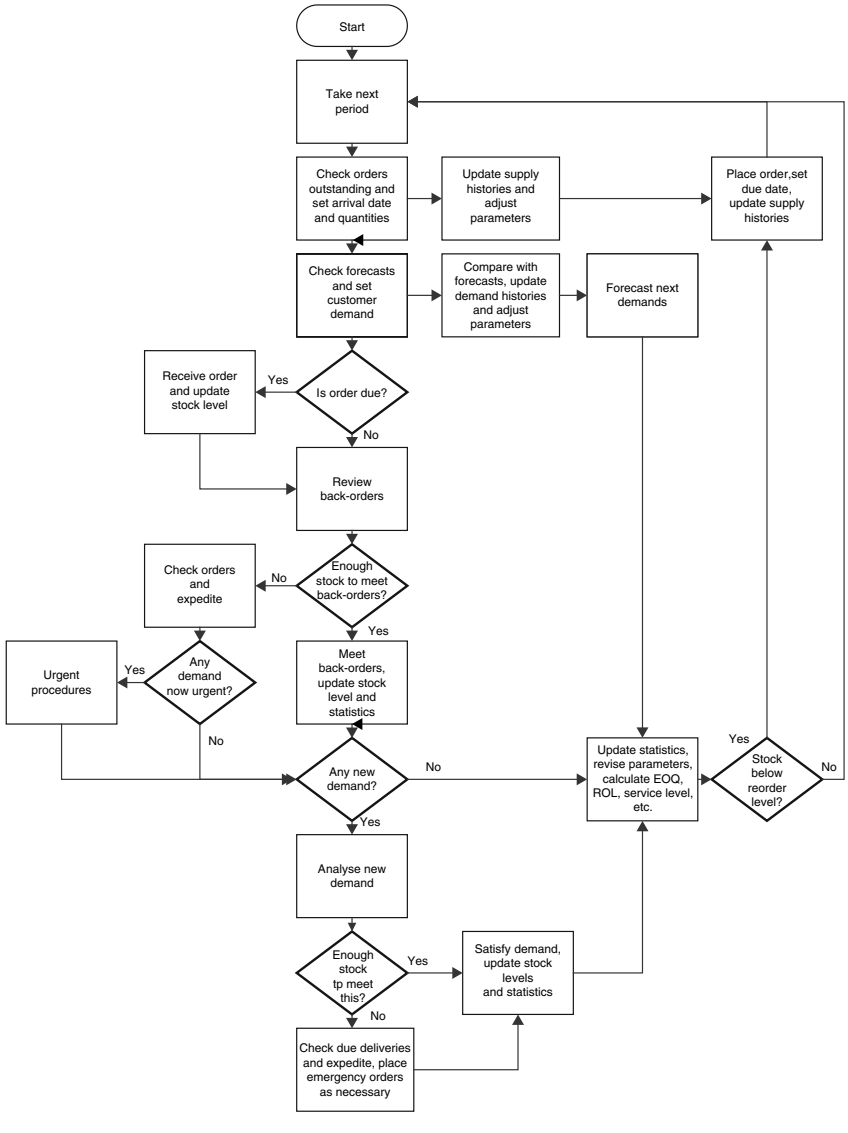


Figure 8.14 Adding some details to the simulation model

start by checking the supply history and setting the pattern of orders due, and then checking the demand history and setting the patterns of customer demands. These, together with associated costs and any relevant details, set the current state of stocks, supplies and demands. Now we can look at the specific questions concerning stocks this period, and in particular how well the supply matches demand. We start by seeing whether a delivery is due and has to be added to stock. Then we look at the demand and see if there are any back-orders outstanding. If possible, these are satisfied, but if there are still shortages expediting and urgent orders are considered. When back-orders have been satisfied, the model moves on to new demand. If there is enough stock, this demand is satisfied, and otherwise expediting and emergency orders are considered. After this, statistics for performance are updated including recalculation of order quantity, reorder level and service level. Then a check is made to see if the stock has fallen below the reorder level, and if it has another order is placed and the model moves on to the next period.

This is clearly one view of part of an inventory system and is at an early stage of development. When this logic was put into a simulation programme some of the results are given in Figure 8.15. Now we can continue expanding the basic model until it accurately describes the real system.

SIMULATION SYSTEMS

SYSTEM: Inventory Management Originator:
Created on 01.01.11 Modified on 01.01.11 This run at 0000 on 01.01.11

DATA REVIEW

Order quantity	Automatic
Reorder level	Automatic
Forecasting method	Exponential smoothing
Forecasting parameters	Optimized
Number of stock items	1

Stock item .1

Name	Component		
Unit cost	£20	Demand history	Defined
Reorder cost	£75	Demand distribution	Normal
Holding cost	£1	Lowest demand	100
Fixed holding cost	£0	Highest demand	500
Shortage cost	£50	Lead time distribution	Uniform
Fixed shortage cost	£0	Shortest lead time	3
Back-orders	0	Longest lead time	5
Urgent orders	No	Service level	93%
Emergency orders	No		

Period	Week		
Number of periods	13	Number of runs	100
Random number seed	Random	Analysis	None
Reports	Transaction (#1,1),	Summary (1)	

Figure 8.15 Part of the output from an inventory simulation

INITIAL VALUES

Opening stock	1,500	Mean lead time	4
Mean demand	300	Lead time demand	1,200
Reorder quantity	212	Safety stock	83
Outstanding orders	Defined	Reorder level	1,283

SIMULATION RESULTS

TRANSACTION REPORT(1) – first run

Week	Opening	Demand	Closing	Shortage	Order	Arrival
1	1,500	331	1,169			
2	1,169	372	797		210 #1	
3	797	229	568		210 #2	
4	568	205	363		210 #3	
5	363	397	0	34	250 #4	
6	0	227	0	227	260 #5	
7	215	0	0	215	270 #6	210 #1
8	210	326	0	116	280#7	210#2
9	210	329	0	119	290#8	210#3
10	210	336	0	126	300#9	250#4
11	250	295	45		310#10	260#5
12	305	280	25		310#11	270#6
13	295	263	32		310#12	280#7

SUMMARY RESULTS(1)

Number of periods	13		
Number of runs	100		
Demand	296.7		
Lead time	5.13 weeks		
Stock level	396.2 units	Holding cost	396.20
Shortages	812.3 units	Shortage cost	40,615.00
		Reorder cost	787.50
		Variable cost	41,798.70
Deliveries	1681 units	Fixed cost	33,620.00
		Total cost	**75,418.70**

Figure 8.15 (*continued*)

Summary

Simulation gives a dynamic representation of real operations. It is based on a simulation model that replicates the real operations over some time, including random effects to get typical results. It is particularly useful for inventory systems where interactions and related effects can be too complex for analytical models.

Review questions

8.16 What is 'simulation'?

8.17 Why is simulation useful for inventory management?

Chapter review

- Planned operations can give an important input to inventory management. This chapter describes the design of these plans and their effects on stocks.
- Planning within an organization is done at different levels, which we described as strategic, tactical and operational. Decisions about stock are made at all levels and they are linked to other decisions within an organization at all levels.
- There must be enough capacity for stocks to meet forecast demand, and they should not form a bottleneck that constrains other operations. There is a general procedure for capacity planning which finds the mismatch between supply and demand, designs plans to overcome this, and implements the best.
- Strategic planning is followed by more detailed tactical plans. We described these as aggregate plans (typically showing monthly production of families of products) and master schedules (typically showing weekly production of each product).
- The main aim of aggregate planning is to meet the requirements of higher plans while using resources efficiently. There are many ways of designing aggregate plans, ranging from intuition to mathematical models. Stocks play an important role in the design of these plans.
- The master schedule disaggregates the aggregate plan to give a detailed timetable for making each product, typically each a week. Designing a master schedules is similar in principle to designing an aggregate plan, and the same methods can be used. Because of the subjective decisions and timing, simpler methods are more common.
- Short-term scheduling designs detailed timetables for all the resources in an organization. Because of the difficulty, timing, flexibility and subjective decisions, this is usually done by intuitive methods and heuristics, often based on scheduling rules.
- A control mechanism is needed to compare actual schedules with plans, and make necessary adjustments
- Simulation gives a dynamic representation of real operations. It replicates the real operations over some time, including random effects to get typical results. This is particularly useful for inventory systems where interactions and related effects can be too complex for analytical models.

Project

The aim of this project is to build a simulation model of an inventory system. For this you have to design a model that shows the detailed steps of an inventory system and then find a way of doing the calculations. In the chapter we demonstrated some results using a spreadsheet, and this is certainly one option. Another is

to write the procedures in a suitable programming language. The best option, though, is to use a specialized simulation program that can make building and running the model very easy. There are many programs available, so see which ones are available and try comparing them. Run your simulation for a range of situations and see how these affect the results.

Problems

8.1 Hoskins Coaches plan their capacity in terms of 'coach-days'. They classify business according to 'full day', which are long distance journeys, or 'half day' which are shorter runs. Forecasts show expected annual demand for next year to be 275,000 full day passengers and 750,000 half day passengers. The company has 41 coaches, each with an effective capacity of 40 passengers a day for 300 days a year. Break-downs and other unexpected problems reduce availability to 90 per cent. The company employs 86 drivers who work a nominal 220 days a year, but illness, training and other absences reduce their availability to 85 per cent. If there is a shortage of coaches, the company can buy extra ones for £110,000 or hire them for £100 a day. If there is a shortage of drivers, they can recruit extra ones at a total cost of £20,000 a year, or hire them from an agency for £110 a day. How could the company approach its capacity planning?

8.2 The fixed cost of a process is €110,000, and the operators can increase capacity by using more machines at a cost of €55,000 each. The total output of the operation, measured in some consistent units, is:

Machines	1	2	3	4	5	6	7	8
Output	55	125	230	310	375	435	460	470

How would the operators get the lowest unit cost?

8.3 Sit-Yuen Foo has produced monthly forecasts of demand for a family of products as shown below. At the end of each month he reviews performance, and assigns a notional holding cost of $30 for every unit held in stock. If there are shortages, 20 per cent of orders are lost with a cost of $300 a unit, and the rest are met by back-orders, with a cost of $75 a unit. Each time the production rate changes there is a cost of $22,500. Designed capacity is 600 units a month, but utilization seldom reaches 80 per cent. How can Sit-Yuen Foo design an aggregate plan for the products?

Month	1	2	3	4	5	6	7	8	9	10
Aggregate demand	460	420	390	450	540	375	240	150	100	210

8.4 The aggregate demand for a family of products for the next five months is 380, 240, 540, 400 and 280 units. Normal capacity is 300 units a month, and

this can be increased by up to 20 units a month by working overtime, with sub-contractors able to handle any amount of production. The unit cost is £200 for normal capacity, £250 for overtime and £280 from sub-contractors. It costs £30 to stock a unit for a month, while back-orders have a penalty cost of £200 a month. How can you use a spreadsheet to design an aggregate plan for the products?

8.5 A bicycle manufacturer produces two bicycles, one for ladies and one for men. The aggregate plan has 8,000 bicycles made next month, and 6,400 the month after. Current stocks are 500 men's and 300 ladies' bikes, and the factory has a capacity of 2,200 bicycles a week. Men's bicycles usually account for 60 per cent of sales, and actual orders have been received for the following deliveries.

Week	1	2	3	4	5	6
Men's bikes	1,400	1,200	1,000	700	300	0
Ladies' bikes	2,000	800	400	100	0	0

Design a master schedule for the 8 weeks.

8.6 The aggregate plan of Hendershaw Pipeworks has 18,000, 15,000 and 15,000 metres of pipework made in the next three months. They want a master schedule for two products, Copper A and Stainless. Current stocks are 700 metres of Copper A and 500 metres of Stainless, and the factory has a capacity of 4,500 metres a week. Sales of Copper A are usually twice as high as sales of Stainless, and the following orders have already been received:

Week	1	2	3	5	6	7	8
Copper A	2,100	1,800	1,600	1,100	800	200	0
Stainless	3,000	1,400	700	400	100	0	0

Design a master schedule for the next 12 weeks.

8.7 Eight batches of items are to be scheduled on a piece of equipment. Each batch fully occupies the equipment for the following number of hours.

Batches	1	2	3	4	5	6	7	8
Duration (hours)	12	5	4	16	2	10	7	8
Due date (hours)	25	30	14	22	56	23	49	33

Use different rules to find alternative schedules for these batches.

8.8 The North Bay Club is looking at the number of part-time bar staff it employs. Over the past three weeks it has counted the following average numbers of customers per hour wanting service at the bar.

Week	Mon	Tue	Wed	Thu	Fri	Sat	Sun
1	65	58	76	122	233	287	46
2	83	77	52	121	155	306	51
3	54	91	94	135	308	176	54

Each of the bar staff can serve about 40 customers an hour. Standard work conditions allow each person a minimum of two nights work a week and a maximum of three nights, and that any nights worked are consecutive. How would you start scheduling the staff?

Discussion questions

8.1 Is it always possible to find the capacity of a process? How, for example, could you find the capacity of a holiday beach, a national park or a shipping lane? Can you give examples where it is difficult to find a capacity, and say how these difficulties are overcome?

8.2 'Planning is only done by manufacturers who keep stocks of their goods. Services do not need this planning, as they are forced to match production to customer demand.' Do you think this is true? What kind of planning would you expect to see in a hospital or firm of management consultants?

8.3 All plans in an organization are derived from the long-term forecasts and capacity plans. But these are notoriously inaccurate. Does this mean that all planning is based on little more than informed guesswork?

8.4 Why do people say that planning is difficult? Plans simply give a timetable for operations – surely it cannot be so complicated to design these?

8.5 Intuitive methods of planning are little more than guesses. The best way to approach planning is to use mathematical analyses that guarantee optimal solutions. Do you agree with this?

8.6 Spreadsheets are widely used for different aspects of planning. Why and how are they used?

8.7 'Simulation can always identify the best plans.' What does this mean, and is it true?

References and further reading

Banks, J., Carson, J.S., Nelson, B.L. and Nicol, D.M. (2001) *Discrete Event Simulation*, Englewood Cliffs, NJ: Prentice-Hall.

Brandimante, P. and Villa, A. (1999) *Modeling Manufacturing Systems*, Berlin: Springer-Verlag.

Brooks, R., Robinson, S. and Lewis, C. (2001) *Simulation and Inventory Control*, Basingstoke: Palgrave.

Browne, J., Harhen, J. and Shivnan, J. (1996) *Production Management Systems* (2nd edition), Reading, MA: Addison-Wesley.
Comel, J.G. and Edson, N.W. (1995) *Gaining Control: Capacity Management and Scheduling*, New York: John Wiley.
Evans, J.R. and Olson, D.L. (2001) *Introduction to Simulation and Risk Analysis*, Englewood Cliffs, NJ: Prentice Hall.
Fogarty, D.W., Blackstone, J.H. and Hoffman, T.R. (1991) *Production and Inventory Management* (2nd edition), Cincinnati, Ohio: South-Western.
Hess, R. (1997) *Managerial Spreadsheet Modelling and Analysis*, Chicago: Irwin.
Klammer, T.P. and Klammer, T. (1996) *Capacity Management and Improvement*, Homewood, IL: Irwin.
Morton, T.E. and Penticto, D.W. (1993) *Heuristic Scheduling Systems*, New York: John Wiley.
Proud, J.F. (1999) *Master Scheduling* (3rd edition), New York: John Wiley.
Rubinstein, R. and Melamed, B. (1998) *Modern Simulation and Modelling*, New York: John Wiley.
Vollman, T.E., Berry, W.L. and Whybark, D.C. (1996) *Manufacturing Planning and Control Systems* (4th edition), New York: McGraw-Hill.

Part IV

Methods for Dependent Demand

9

Material Requirements Planning

Aims of the chapter

This chapter introduces material requirements planning (MRP), which is based on the idea that we can use planned production to find the demand for materials. MRP starts with the master schedule and then uses a bill of materials to transform this into a timetable of required materials. We can use this to schedule orders sent to suppliers and related internal operations.

After reading this chapter you should be able to do the following:

- discuss some of the limitations of independent demand methods;
- describe the characteristic approach of material requirements planning (MRP);
- use MRP to timetable orders and operations;
- outline the benefits and problems of using MRP;
- adjust MRP schedules with batching rules;
- add feedback to check for capacity and other problems;
- discuss some extensions to MRP, including MRP II, DRP and ERP.

This chapter emphasizes:

- *dependent demand methods*, where the demand for materials is found from production plans;
- *material requirements planning*, which explodes a master schedule to give a timetable for supplying materials.
- *extensions* of the basic MRP method.

Limitations of independent demand methods

Methods for independent demand are very widely used. Since they were introduced in the 1920s they have proved a valuable and flexible tool for management. Nonetheless, there are some circumstances in which they do not perform well.

One problem area is spare parts for equipment or other slow moving items. Here the demand is so low that forecasts based on historical demand are close to zero, and the item is never in stock. Special models have been built for these situations (which we mentioned in Chapter 5) but the best policy is often a simple rule along the lines of 'order a replacement unit whenever one is used'.

Another problem comes in batch production, when units pass through a process in discrete batches. Imagine the operations in one part of a process working on a batch of one item; when they have finished work on this batch, they move on to a batch of another item. But the two batches need different materials. So the demand for some materials suddenly falls from a high level down to zero, while the demand for other materials goes from zero up to some high level. You can see the problem this creates with forecasting demand from historical figures. When an item is not currently being made, the forecasts for its materials fall to zero, none are kept in stock, and they are not available when a new batch is started.

Although this seems like a problem with forecasting, it is really a more fundamental problem with the basic assumption of independent demand. The number of units of an item made determines the demands for all components and raw materials needed to make it. So these demands are not independent, but are related through the production plan. We now have *dependent demands*. By definition, independent demand methods work when demands are independent, but we must question whether they can really work for dependent demands.

In practice, the answer is that independent demand methods often work very well for dependent demands. Even if the basic models need some adjustment, they can easily be modified to give useful results in a variety of different conditions. But having recognized a weakness in the approach, you might wonder if there are other problems. In certain circumstances, independent demand methods do perform less well and they need – at least – some adjustment. The following list gives some common problems:

- They assume that demands for items are independent. In reality, there are many circumstances – not just raw materials for production – where there is some mechanism linking demands.
- They assume that demand is relatively stable, or follows some other pattern that can be accurately forecast from historical figures. But historical figures may be irrelevant to the future, or it may be difficult to identify a realistic pattern.
- They assume that stock should always be available, even during periods when there is no expected demand.
- They forecast demand from historical values even when future demand can be found with certainty from production plans or some other source.
- They are not really suitable for forward planning (especially over the longer term) as all the variables, calculations and decisions are based on fairly recent historical values.
- Variables, such as the lead time, demand and costs, are assumed to be outside the organization's control. In reality, they can be varied by negotiation, using different suppliers, expediting, emergency procedures, and so on.

- Cost are assumed to be fixed, even though the reorder cost, for example, can vary with the supplier, distance to travel, procedure for ordering, people involved, etc.
- Costs rely on accounting conventions and estimates. This is obvious for shortage costs, but is also true of holding and reorder costs.
- Variables such as lead time are assumed to follow known distributions, but these can be difficult to identify.
- Even if a high service level is set for all materials, a product using many materials is unlikely to find all of them in stock at the same time. When a service level of 99 per cent is used for parts, a product that is made from 100 parts has a probability of $0.99^{100} = 0.37$ of finding all of them in stock.
- Informal procedures affect operations – perhaps exerting pressure to use expediting and emergency procedures, or keeping unofficial safety stocks.
- The models accept the status quo and give no incentive or mechanism to find improvements.

There can also be weaknesses in specific situations. When, for example, an organization looks for a new supplier, there may be an intricate bidding process, long negotiations, special design modifications, or some other complication. Then the reorder cost is very high, this raises the economic order quantity, and orders may be so large that they exceed the lifetime requirement, or the item becomes obsolete before all units are used.

We can conclude that independent demand methods are very good at dealing with many problems of inventory control, but they do have weaknesses. This should encourage us to look for another approach. One of the most widely used is *material requirements planning* (MRP). This is a dependent demand method that is useful when demands for materials are related through a production plan.

Summary

Despite their widespread use, independent demand methods have some weaknesses. We can modify the basic models, but in some circumstances – such as batch production – an alternative approach might be better. Material requirements planning gives a widely used alternative for dependent demand.

Review questions

9.1 What is the essential difference between independent demand and dependent demand methods of inventory control?

9.2 Independent demand methods cannot be used for batch production or sporadic demand. Do you think this is true?

Approach of material requirements planning

Overall aims

In its basic form material requirements planning (MRP) takes a master schedule and uses this to design a detailed timetable for ordering materials. The master schedule shows the number of units of an item made, typically in each week. MRP takes this, together with a list of the materials needed for each unit, and develops a timetable for the supply of materials. These materials are either bought in or made internally, so the main outputs from MRP are:

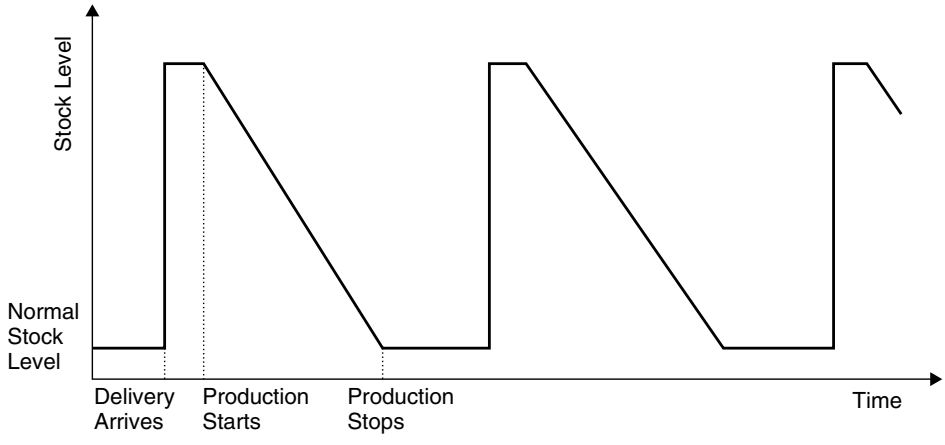
- timetables to show when materials are needed;
- timetables to show when bought-in materials should be ordered;
- timetables for operations needed to make materials internally.

- *Material requirements planning* 'explodes' the master schedule to plan the supply of materials.
- It gives *timetables* for making and ordering materials, to make sure that they are available when needed.

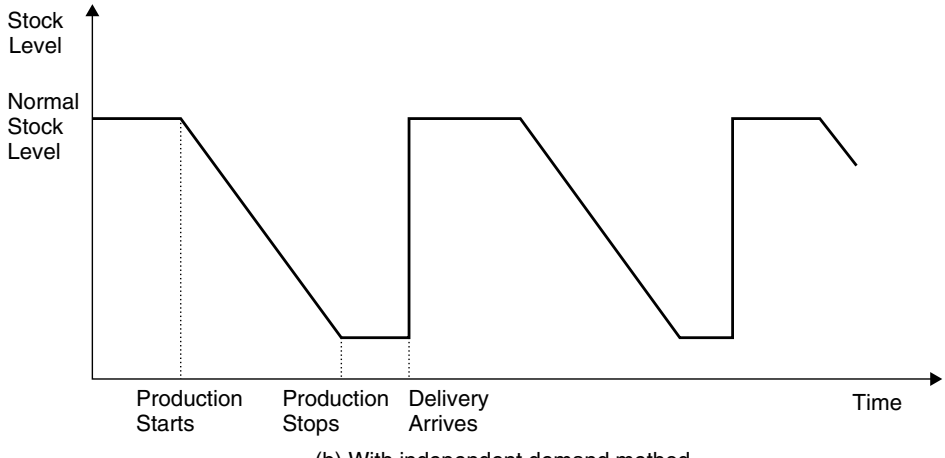
This approach matches the supply of materials directly to the known demand, so there is only enough stock to satisfy this demand. In contrast, the alternative independent demand methods keep stocks of materials that are high enough to cover any likely demand. An analogy is the way that a chef plans the ingredients needed to cook a week's meals. The MRP approach looks at the menus for each day, uses this to find the ingredients needed, and then orders these to be delivered in time for their use. Independent demand methods see what ingredients were used in previous weeks and orders enough of everything make sure that stocks are high enough to cover likely demand in the future.

An important point is that the two approaches give different patterns of material stocks. With MRP stocks are generally low, but rise as deliveries are made just before production starts. Stock is then used during production and the amount held declines until it returns to a normal, low level. This pattern is shown in Figure 9.1(a). With independent demand methods stocks are not related to production plans, so higher levels are kept in case they are needed. These are reduced during production, but are replenished as soon as possible, to give the pattern shown in Figure 9.1(b). An obvious benefit of MRP, then, is a lower average stock level.

MRP was originally developed for manufacturing industry. Although it is now used more widely we will, for convenience, stick to the original vocabulary which talks of stocks of material (parts and components) being delivered to make products.



(a) With dependent demand method, MRP



(b) With independent demand method

Figure 9.1 Comparison of stock levels

MRP procedure

MRP needs a lot of information, with three main sources being the master schedule, the bill of materials and inventory records. We know that the master schedule shows the number of units of a product to be made in each period, and the inventory records show the current state of the stocks. The *bill of materials*, or parts list, is an ordered list of all the materials needed to make a product, and also the order in which the materials are used.

Suppose that a company assembles desks from a top and four legs. A common format for the bill of materials is shown in Figure 9.2. As you can see, every item has a 'level' that shows where it fits into the process, and figures in brackets show the numbers needed to make each unit. The finished product is level 0; level 1 items are used directly to make the level 0 item.

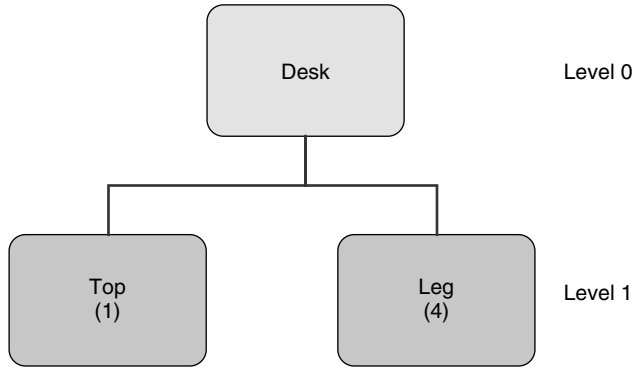


Figure 9.2 Initial bill of materials for a desk

You will often see bills of materials simplified into a structured list, and for the desk we could write:

- Level 0
  Desk
  - Level 1
    Top (1)
    Leg (4)

Now suppose that the master schedule shows that six desks are to be made in the week beginning 21 June. MRP can 'explode' the master schedule using the bill of materials to give details of the materials needed. The six desks need six tops and 24 legs, and these have to be in stock ready for assembly by 21 June. These are actually the *gross requirements*. We may not have to order all of these, as the inventory records may show some already in stock, or with outstanding orders that are due to arrive shortly. If we subtract these from the gross requirements we get the *net requirements* for materials, which are the amounts we have to order. We need 24 desk legs by 21 June, but if we already have four in stock and an order of 12 that is due to arrive shortly, our net requirement is for $24 - 4 - 12 = 8$.

For each material:

- Gross requirements = number of units made × amount of material for each unit
- Net requirements = gross requirements − current stock − stock on order

Now we know the quantities to order, and when these orders should arrive. The next step is to find the time to place the orders. For this we need the lead times and we place orders this lead time before the materials are actually needed. If we buy the legs from suppliers who give a lead time of three weeks, we need to place orders at the beginning of June. This is called *time shifting*. If we have no stocks

of desk tops but buy them from suppliers who give a lead time of two weeks, we have to order 6 by the 7 June. This gives the following timetable:

1 June:	order 8 legs
7 June:	order 6 tops
21 June:	start assembly of desks

Before finalizing these orders we have to consider any other relevant information, such as minimum order sizes, discounts, minimum stock levels, variation in lead time, and so on. When we take all of this into account we can get detailed timetable for orders. This procedure is shown in Figure 9.3.

It is easiest to see these calculations in a table, and Figure 9.4 shows them in a spreadsheet. Here rows 4 to 9 show the schedule for level 0 items, which is the

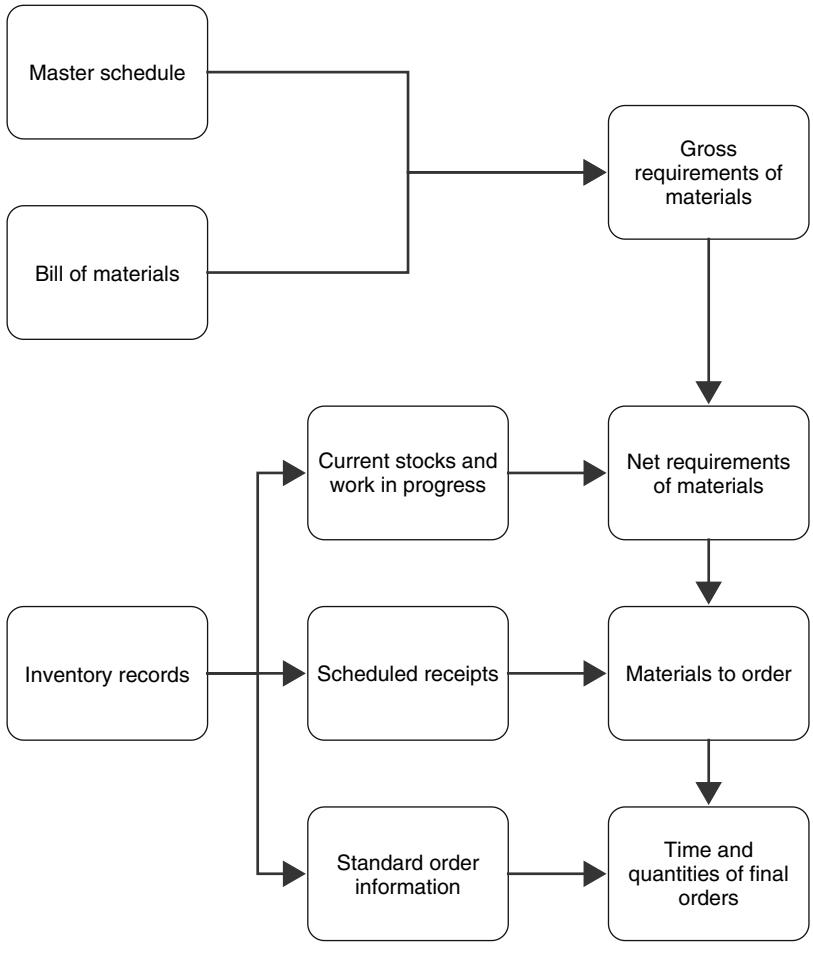


Figure 9.3 Summary of MRP calculations

	A	B	C	D	E	F	G
1	**Material requirements planning**						
2							
3		**24th May**	**1st June**	**7th June**	**14th June**	**21st June**	**28th June**
4	**Level 0 - desks**						
5	Gross requirements					6	
6	Opening stock						
7	Scheduled receipts						
8	Net requirements					6	
9	Start assembly					6	
10							
11	**Level 1 - tops (1 per desk, lead time 2 weeks)**						
12	Gross requirements					6	
13	Opening stock						
14	Scheduled receipts						
15	Net requirements					6	
16	Place order			6			
17							
18	**Level 1 - legs (4 per desk, lead time 3 weeks)**						
19	Gross requirements					24	
20	Opening stock	4	4	16	16	16	
21	Scheduled receipts			12		8	
22	Net requirements					8	
23	Place order		8				

Figure 9.4 Spreadsheet of the MRP calculations

desk. The gross requirement in row 5 (for six units starting on 21 June) comes from the master schedule. This is translated into net requirements in row 8, and into assembly plans in row 9. Rows 11 to 16 show the schedule for tops and rows 18 to 23 show the schedule for legs. The gross requirements for these level 1 items are set by the net requirements for level 0 items. Then subtracting the opening stock and scheduled receipts gives the net requirements. Moving backwards by the lead time gives the time to place orders, and these are translated into scheduled receipts.

We have only considered two levels of the bill of materials, but if we look in more detail we will probably find many other levels. We could make each top from a wood kit and hardware, with the wood kit having four oak planks, two pine insets and four veneers: the oak planks might consist of two planed 3 cm by 30 cm, and so on. Part of this more detailed bill of materials is shown in Figure 9.5, with level 2 materials used to make the level 1 items, level 3 materials used to make level 2 items, and so on. The full bill of materials keeps going down through different levels until it reaches materials that the organization always buys in from suppliers. By this time, there might be hundreds or even thousands of different materials.

We can summarize the overall MRP procedure in the following steps:

1. Use the master schedule to find the gross requirements of level 0 items.
2. Subtract any stock on hand and orders arriving to give the net requirements for level 0 items. Then schedule production, with starting times to meet these net requirements.

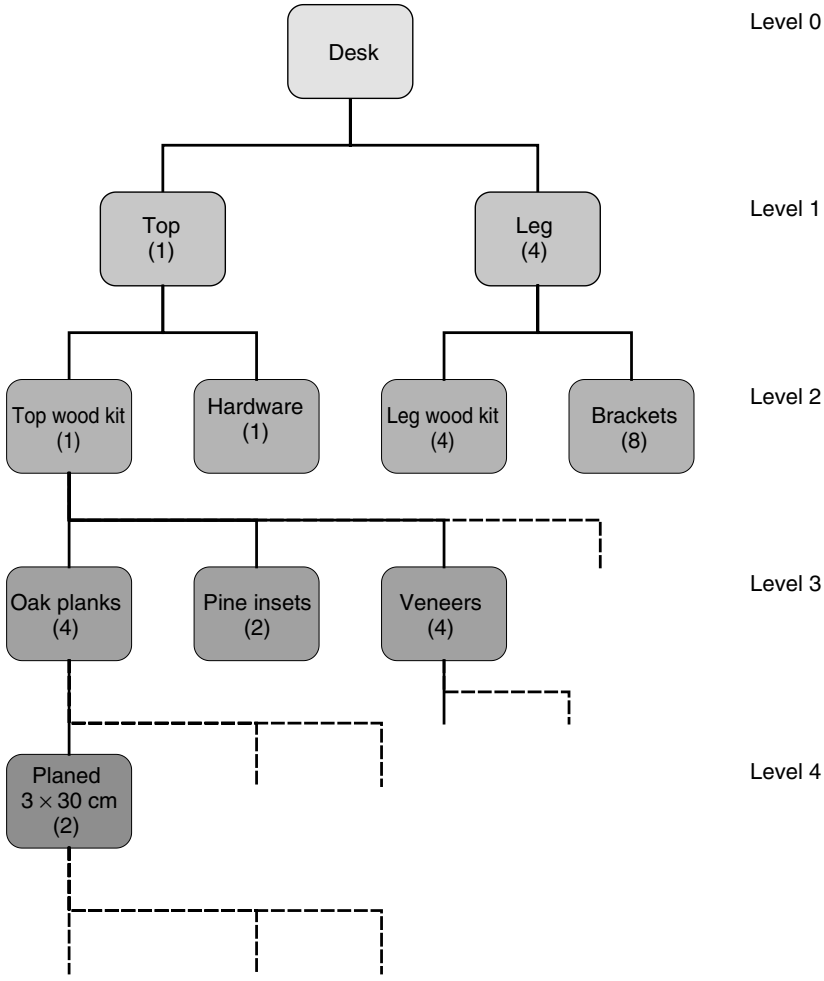


Figure 9.5 Expanding a bill of materials for a desk

3. If there are more levels of materials, use the bill of materials to translate the assembly or orders from the last level into gross requirements for this level. If there are no more levels go to step 5.
4. Take each material in turn and:
   - subtract the stock on hand and scheduled deliveries to find the net requirements, which are the amounts to order;
   - use the lead time and any other relevant information to give the timing of these orders;
   - go back to step 3.
5. Finalize the timetable for orders and production, adding any specific adjustments.

Worked example

MarkMobil Ltd. assembles supermarket trolleys from a main body and four wheel assemblies. They make the body themselves from a body kit and two handle kits. The lead times are one week for assembling trolleys, three weeks for buying wheels, one week for assembling the body, three weeks for buying the body kit and one week for the handle kit. MarkMobil receive orders for 100 trolleys to be delivered in week 8 of a production period, and 200 trolleys in week 10. It has stocks of 20 complete trolleys, 110 bodies and 200 wheels, but no body kits (which must be bought in batches of 200) or handles (which must be bought in batches of 400). Design a timetable for production of the trolleys.

Solution

The bill of materials for this problem is shown in Figure 9.6. The calculations are in the spreadsheet of Figure 9.7, and you can follow the standard procedure described above.

1. Level 0 items are the finished products, so we start by looking at the demand for trolleys in rows 4 to 9, with the gross requirements in row 5.
2. Subtracting the stocks of finished trolleys from the gross requirements gives the net requirements in row 8. Assembly takes one week, so we get the start times in row 9, and these are translated into scheduled receipts in row 7.
3. Now we use the assembly plans for Level 0 items (trolleys) to find the gross requirements for level 1 items (bodies and wheels). In week 7 the company starts assembly of 80 trolleys, and the bill of materials translates this into gross requirements for 80 bodies and 320 wheels. Similarly, in week 9 the company

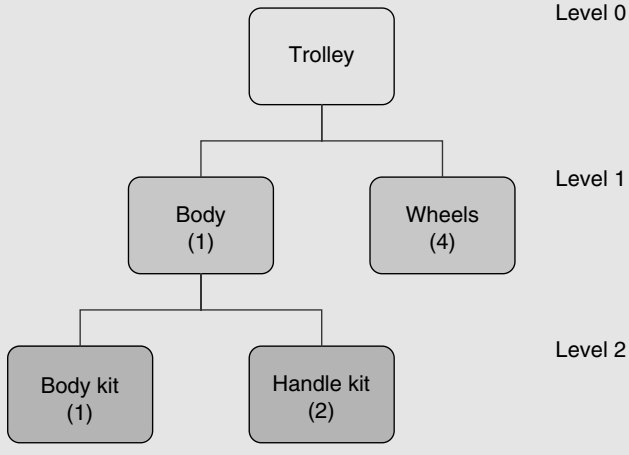


Figure 9.6 Bill of materials for a trolley

	A	B	C	D	E	F	G	H
1	**Material requirements planning**							
2								
3	Week	4	5	6	7	8	9	10
4	**Level 0 - trolley (assembly time 1 week)**							
5	Gross requirements					100		200
6	Opening stock	20	20	20	20	20		
7	Scheduled receipts					80		200
8	Net requirements					80		200
9	Start assembly				80		200	
10								
11	**Level 1 - body (1 per trolley, assembly time 1 week)**							
12	Gross requirements				80		200	
13	Opening stock	110	110	110	110	30	30	
14	Scheduled receipts						170	
15	Net requirements						170	
16	Start assembly					170		
17								
18	**Level 1 - wheels (4 per trolley, lead time 3 weeks)**							
19	Gross requirements				320		800	
20	Opening stock	200	200	200	200			
21	Scheduled receipts				120		800	
22	Net requirements				120		800	
23	Place order	120		800				
24								
25	**Level 2 - body kit (1 per body, lead time 3 weeks, order size 200)**							
26	Gross requirements					170		
27	Opening stock							
28	Scheduled receipts					200	30	30
29	Net requirements					170		
30	Place order		200					
31								
32	**Level 2 - handle kit (2 per body, lead time 1 week, order size 400)**							
33	Gross requirements					340		
34	Opening stock						60	60
35	Scheduled receipts					400		
36	Net requirements					340		
37	Place order							

Figure 9.7 Spreadsheet of calculations for the worked example

starts assembling 200 trolleys, and this translates into gross requirements for 200 bodies and 800 wheels.

4. Subtracting stock on hand from these gross requirements gives the net requirements. To make sure that the parts arrive on time they must be ordered the lead time in advance (i.e. one week for bodies and three weeks for wheels). These calculations are shown in rows 11 to 16 for bodies and 18 to 23 for wheels.

5. Now we repeat this procedure for level 2 items, where the gross requirements for body kits and handle kits comes from the assembly plans for bodies. These calculations are shown in rows 25 to 30 for body kits and 32 to 37 for handle kits. The only difference here are the fixed order sizes, with any unused deliveries kept in stock.

6. There are no more levels of materials, and we know of no special conditions, so we can finalize the timetable of events as:
   week 4: order 120 wheels
   week 5: order 200 body kits
   week 6: order 800 wheels
   week 7: start assembling 80 trolleys, receive 120 wheels, order 400 handle kits
   week 8: finish 80 trolleys, start assembling 170 bodies, receive 200 body kits, receive 400 handle kits
   week 9: start assembling 200 trolleys, receive 170 bodies, receive 800 wheels
   week 10: finish assembly of 200 trolleys

Worked example

The master schedule of Parsival Fabricators is designed to meet demands of 90 units of a product in week 7 of a cycle, 120 units in week 8 and 80 units in week 11. They currently have 20 units of the product in stock, but the company always keeps 10 units in reserve to cover emergency orders. Each unit of the product takes two weeks to assemble from 2 units of component B and 3 units of component C. Each unit of component B is made in one week from 1 unit of part D and three units of part E. Component C is assembled in 2 weeks from two units of material F. Lead times for D, E and F are 1, 2 and 3 weeks respectively. Current stocks are 100 units of B, 200 of C, 80 of D, 720 of E and 200 of F. Parsival keep minimum stocks of 40 units of D, 200 of E and 100 of F. The minimum order size for E is 600 units, while F can only be ordered in discrete batches of 200 units. Outstanding orders for 20 units of D will arrive in week 4, 600 units of E will arrive in week 5, 400 units of F will arrive in week 4, and 40 units of C will arrive in week 6. Design a timetable of activities for the company.

Solution

Even a simple MRP problem needs a lot of calculation and spreadsheets soon become very complicated. A better option is to use specialized software. Figure 9.8 shows the printout for this problem from a basic program.

The program reviews the bill of materials, and then starts the analysis at level 0, with production of the final product, A. Parsival keep a minimum stock of 10 units of A, and this reserve stock must be set aside when calculating the net requirements.

Then the program moves on to level 1 materials and expands the assembly plan for A into gross requirements for components B and C. The 80 units of A assembled in week 5 are expanded into gross requirements of 160 units of component B and 240 units of component C. The 120 units of A assembled in week 6 are expanded into gross requirements of 240 units of B and 360 units

```
Material Requirements Planning


Parsival Fabricators
Analysis  originated  at  00.00  on  11.11.11  by    XXXX  XXXX


Period:             weeks
Number of weeks:  13
Analysis for weeks:4 to 11


Bill of Materials

Level 0 End Item Product A
          |
          --------------------------------
          |                |
 Level 1 X Component B                    X Component C
          | 2 units |                       |3    units
          |         |                       |
          ----------------                  |
          |                |                |
Level 2   X  Part D       X  Part E         X Material F
          |  1 unit       |  3 units        | 2 units
          |               |                 |

MRP Analysis - weeks 4 to 11

Level 0 - End Item
  ItemNumber : Part-0                Opening stock: 20
  Description: Product A             Lead time:     2
                                     Safety Stock:  5
                                     Lot Size:      1

Week:                    4    5    6    7    8    9    10   11
 Gross requirements:     0    0    0    90   120  0    0    80
 Stock              :    20   20   20   20   10   10   10   10
 Net requirements   :    0    0    0    80   120  0    0    80
 Receipts           :    0    0    0    80   120  0    0    80
 Requests           :    0    80   120  0    0    80   0    0

Level 1 - Comp 1
 Item number: Part-1                Opening stock: 100
 Description: Component B           Leadtime:      1
 Number     : 2                     Safety stock:  0
 For part   : Part-0                Lot size:      1

Week:                      4   5    6    7   8    9    10   11
 Grossrequirements  :      0   160  240  0   0    160  0    0
 Stock              :    100   100  0    0   0    0    0    0
 Net requirements   :      0   60   240  0   0    160  0    0
 Receipts           :      0   60   240  0   0    160  0    0
 Requests           :     60   240  0    0   160  0    0    0

Level 2 - Comp 1-1

 Item number: Part-2                Opening stock: 80
 Description: Part D                Leadtime:      1
 Number     : 1                     Safety stock:  40
 For part   : Part-1                Lot size:      1

Week:                      4     5    6    7    8     9    10   11
 Gross requirements  :     60    240  0    0    160   0    0    0
 Stock               :     80    40   40   40   40    40   40   40
 Net requirements    :     20    240  0    0    160   0    0    0
 Receipts            :     20    240  0    0    160   0    0    0
 Requests            :     240   0    0    160  0     0    0    0
```

Figure 9.8 Printout from a basic MRP program for the worked example

Level 2 - Comp 1-2

Item Number: Part-3					Opening stock: 720				
Description: Part E					Lead Time: 2				
Number : 3					SafetyStock: 200				
For part : Part-1					Lot Size: 600				
Week:	4	5	6	7	8	9	10	11	
Gross requirements :	180	720	0	0	480	0	0	0	
Stock :	720	540	420	420	420	540	540	540	
Net requirements :	0	380	0	0	260	0	0	0	
Receipts :	0	600	0	0	600	0	0	0	
Requests :	0	0	600	0	0	0	0	0	

Level 1 - Comp 2

Item Number: Part-4				Opening stock: 200				
Description: Component C				Lead Time: 2				
Number : 3				SafetyStock: 0				
For part : Part-0				Lot Size: 1				
Week:	4	5	6	7	8	9	10	11
Gross requirements :	0	240	360	0	0	240	0	0
Stock :	200	200	0	0	0	0	0	0
Net requirements :	0	40	360	0	0	240	0	0
Receipts :	0	40	360	0	0	240	0	0
Requests :	360	0	0	240	0	0	0	0

Level 2 - Comp 2-1

Item Number: Part-5				Opening stock: 200				
Description: Material F				Lead Time: 3				
Number : 2				SafetyStock: 100				
For part : Part-4				Lot Size: 200				
Week:	4	5	6	7	8	9	10	11
Gross requirements :	720	0	0	480	0	0	0	0
Stock :	200	280	280	280	200	200	200	200
Net requirements :	620	0	0	300	0	0	0	0
Receipts :	800	0	0	400	0	0	0	0
Requests :	400	0	0	0	0	0	0	0

Figure 9.8 (*continued*)

of C, and so on. These gross requirements for B and C can be met partly from opening stocks, with the shortfall shown as net requirements. We also have to remember the planned delivery of 40 units of C in week 5.

This schedule for level 1 items is expanded to give the timetable for level 2 items. The gross requirements for parts D and E come from the assembly plans for component B. 60 units of B are started in week 4 and this expands into gross requirements for 60 units of D and 180 units of E, and so on. One complication here is the minimum order size of 300 units of E. In week 5 there is a gross requirement of 720 for E, 340 of which can be met from free stock (keeping the reserve stock of 200). The net requirement is 380, but 600 have to be ordered with the spare 220 added to stock.

Finally, the gross requirements for material F can be found from the assembly plan for component C. 360 units of C are started in week 4 and this expands into a gross requirement of 720 units of F, and so on. Orders must be in discrete

batches of 200 units, so they are rounded to the nearest two hundred above net requirements.

The overall timetable of activities is:

week 4:	start making 60 of B and 360 of C
	place orders for 240 units of D and 400 units of F
	orders arrive for 20 units of D and 800 units of F
week 5:	start making 80 of A and 240 of B
	finish 60 units of B
	orders arrive for 40 units of C, 240 units of D and 600 units of E
week 6:	start making 120 of A
	finish 240 units of B and 360 units of C
	place order for 600 units of E
week 7:	finish making 80 units of A
	start making 240 of C
	place order for 160 units of D
	order arrives for 400 units of F
week 8:	finish 120 units of A
	start making 160 units of B
	orders arrive for 160 units of D and 600 units of E
week 9:	start making 80 units of A
	finish 160 units of B and 240 units of C
week 11:	finish 80 units of A

Summary

Material requirements planning uses a bill of material to explode a master schedule and find the gross requirements for materials needed to support production. Information about current stocks, orders outstanding, reserved stocks, etc. are then used to find order quantities. Lead times are use to time phase these orders so that materials arrive in time for use. The resulting stocks are matched directly to production plans.

Review questions

9.3 What is the basic information needed by MRP?

9.4 How would you find the net requirement for material using MRP?

9.5 The structure of MRP means that it can only be used for manufacturing. Do you think this is true?

Benefits and problems with MRP

Inputs and outputs from MRP

One of the benefits of MRP is the wealth of information that it collects, analyses and distributes. The main inputs are the master schedule, the bill of materials,

inventory records and any other relevant information, and the main output is a timetable for material orders. But with all of this data and analysis, MRP can generate many different kinds of report, including:

- *timetables of operations* needed to achieve the master schedule, particularly times for internal production to make materials;
- *a timetable of orders* for materials from external suppliers;
- *changes to previous orders* – whenever the master schedule is revised, or any other changes are made, the MRP schedules have to be updated. The MRP system can report changes to order quantities, cancelled or changed orders, changes of due dates, and so on;
- *exceptions* – if the MRP system cannot deal automatically with some unusual circumstances it will report these as needing management action. Typical problems include infeasible schedules, late orders, overloaded capacity, excess scrap, requests for non-existent parts, shortage, and so on;
- *performance reports* – which show how well the system is working, including measures for investment in stocks, inventory turnover, costs and number of shortages;
- *planning reports* – which give information for longer-term planning decisions, including capacity requirements and achievable objectives;
- *records of inventory transactions* – keeping accurate records of current stocks, planned positions and monitoring progress.

These give the overall view of an MRP system shown in Figure 9.9.

This ability to generate a lot of reports is a clear benefit of MRP, but some people say that it can also become a problem (and that MRP stands for 'More Reams of Paper'). MRP is typically run every week, and it can swamp managers with poorly prepared reports, analyses they do not want and unimportant details. The implication is that the MRP reports have to be carefully designed so that each manager is only given relevant and appropriate information.

Other benefits

An obvious benefit of MRP is its ability to relate stocks directly to planned operations. Traditional, independent demand methods forecast likely demand for materials, and then hold stocks that are high enough to meet these. Unfortunately, we know that forecasts are usually wrong, and to allow for the errors organizations hold more stocks than they really need, and this inevitably increases costs. MRP avoids these costs by relating the supply of materials directly to demand. The result is much lower stocks and related costs.

You might think that as MRP relates stocks directly to demand, there is no need for safety stocks. In practice, most organizations still add some safety stock, arguing that supplies with MRP might hit the same problems as supplies with

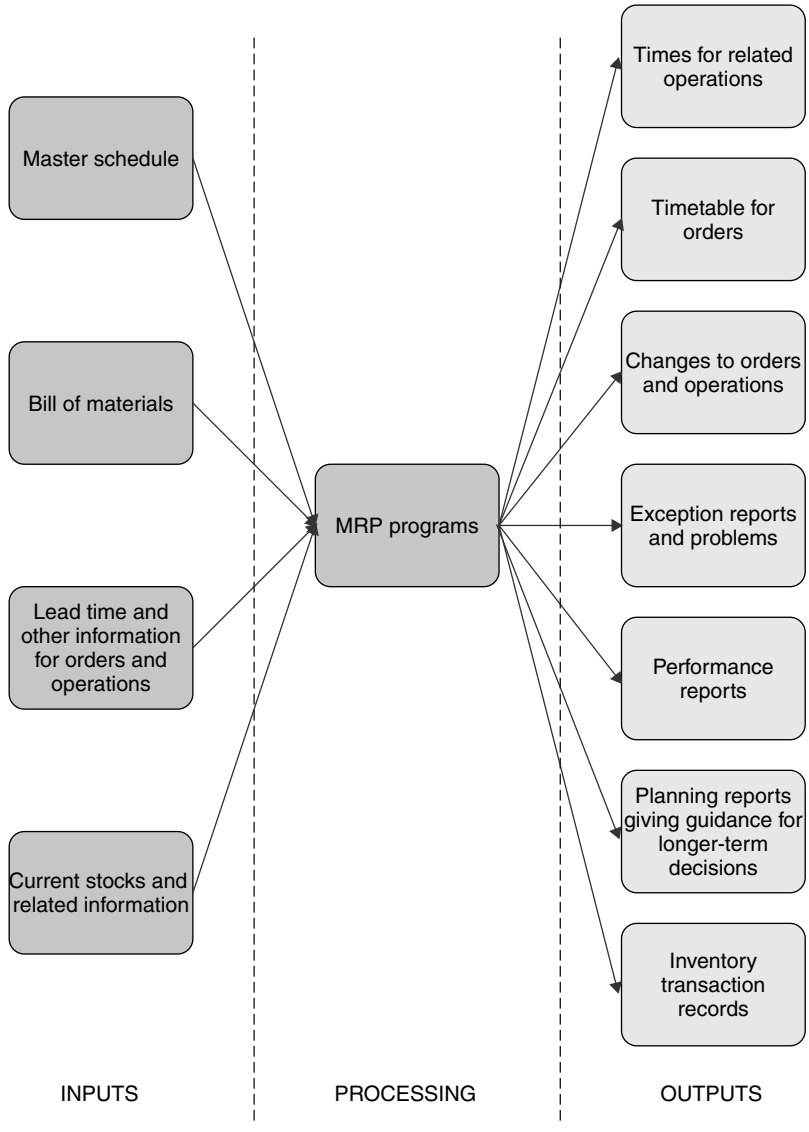


Figure 9.9 Summary view of an MRP system

independent demand methods. Even if MRP tries to eliminate the uncertainty in demand, there are still unforeseen difficulties, and there can still be unexpected variations in supply and lead time. Of course, there is an opposing view which says that organizations can allow for uncertainty by adjusting the lead time given to customers, expediting deliveries, or adjusting production priorities. Then safety stocks are an unnecessary waste of money. This latter view is becoming dominant, and is certainly closer to the underlying objectives and principles of MRP.

The main benefits that come from this direct link between demand for products and the supply of materials, include:

- materials supply is linked directly to known demand;
- lower stock levels, with savings in capital, space, warehousing, etc.;
- higher stock turn-over;
- better customer service – with no delays caused by shortages of materials;
- more reliable and faster delivery times to customers;
- higher utilization of facilities – as materials are always available when needed;
- less time spent on expediting and emergency orders;
- MRP schedules can be used for short-term planning;
- assigns priorities for jobs supplying materials;
- encourages better planning.

This last benefit is particularly interesting. MRP effectively links several aspects of planning, including tactical master schedules and operational scheduling of resources. This puts pressure of organizations to design good plans and then stick to them. The result is better planning. MRP can also give early warning of potential problems and shortages. If necessary, expediting can be used to speed up deliveries, or production plans and priorities can be changed.

Another broad benefit of MRP comes from its detailed analyses that highlight problems which have previously been hidden. For example, an organization might not notice an unreliable supplier if it keeps enough stock to avoid problems. This effectively hides the problem, but increases costs. With MRP, stock levels are much lower and they depend on reliable deliveries. The organization will certainly notice any late deliveries, and they can take steps to improve the supplier's reliability, either by changing the supplier or discussing ways of improving performance.

Problems with MRP

In contrast to these advantages, there are also some problems with MRP. The most obvious is the amount of information and calculation that it needs. MRP starts with a detailed master schedule, so it cannot be used if:

- there is no master schedule;
- the master schedule is not designed far enough in advance;
- the master schedule is inaccurate, not showing what actually happens;
- plans are changed frequently.

Other information needed by MRP includes a bill of materials, information about current stocks, orders outstanding, lead times, and other information about

suppliers. Many organizations simply do not record this information. Others find that their information does not include enough detail, or is in the wrong format for MRP.

Even when an organization seems to have all the information, there can be problems with accuracy. Inventory files, for example, are updated with every transaction, but large numbers of small transactions easily introduce errors. Ordinarily such small errors are not really important, as stocks are high enough to cover problems until the errors are noticed. But MRP does not have these high stocks, so production depends on accurate records. Small errors become important, as stocks are only available for the specified master schedule, and if there are errors in stock records, this schedule cannot be completed.

The sheer volume of information processing can cause difficulties. You can see from the worked examples above that very simple problems soon give a lot of data manipulation, but in reality organizations are likely to work with hundreds or thousands of products and materials. The ideas behind MRP are not particularly new, but they only started to become practical with cheap computing in the 1970s (Orlicky, 1975; Wight, 1974). However, it took another 20 years for reliable, large-scale systems to become widely available. The volume of data processing makes some people say that MRP is not really a method of planning or inventory management, but is a part of the management information system. Its main role is to do a simulation of operations and provide reports and analyses. Other people suggest that this is playing with words, and whatever we call MRP, we should recognize it as a valuable tool.

The information needed by MRP clearly limits its use. A manufacturer is likely to work from a master schedule and can use MRP for raw materials and work in progress: a supermarket is unlikely to have a detailed master schedule and must use some other method. MRP was originally designed for manufacturers, but many other types of organization have adopted it. Universities, airlines, dentists, consultants and many other services know in advance which customers they will serve. Then they can use these known schedules to plan the facilities they need. Package holiday companies, for example, design a schedule of holidays, and they can use this to plan the materials they need (hotel rooms, flights, meals, transfers, etc.). When we talk about materials being delivered for products, remember that – as always – our products are a combination of goods and services.

People often suggest that inflexibility is another serious problem with MRP. Independent demand methods typically hold large stocks of all materials, and these allow production plans to be changed at short notice; with MRP the only materials available are those needed for the specified master schedule, so an organization cannot modify this schedule to allow for changing circumstances. There are simply no materials available to make any other schedule. In practice, an organization can increase flexibility by keeping some safety stocks, or regularly updating of the MRP schedule. The updating can be done in two ways:

- have a complete MRP run every period, typically every week. This is 'regenerative MRP' and has the advantage of regularly taking into account all new information. On the other hand, it has the disadvantage of needing a lot of processing and perhaps giving major changes.

VELVSHER

- update the schedules when there are changes. This is 'net change MRP' which can be a limited run to show the effects of changes. This updating is done more frequently (typically daily) and has the advantage of reducing the amount of processing by only recalculating and reporting changes from previous periods. On the other hand, it has the disadvantage of allowing frequent changes that can make the process seem 'nervous'.

Despite these adjustments to increase flexibility, MRP remains a very rigid and formal system. This formality can be a disadvantage, as organizations sometimes concentrate so much on the plans that they begin to see the meeting of schedules as their highest priority. Everyone's aim is to meet specified targets, and no decisions are made outside the system. In practice, many organizations have informal systems working alongside the formal ones to allow flexibility and make sure that things get done. MRP discourages these informal systems, and this may be seen as a weakness (but, conversely, some people consider it a strength).

Another disadvantage comes with the assumptions that materials are made in the 'bottom-up' order described by the bill of materials. In other words, high-level materials are made before lower level ones. In practice such schedules could give very inefficient workloads, particularly in a complex environment like a job shop. To avoid this, most organizations make parts in an order that is very different from the bill of materials, and improve productivity and efficiency by taking a longer-term view.

One other disadvantage that we should mention in passing is the order pattern suggested by MRP. This may give small frequent orders, which are inconvenient and expensive to administer. We will describe a way to avoid this in the next section.

We can summarize some of the disadvantages of MRP as follows:

- reduced flexibility to deal with changes;
- needs a lot of detailed and reliable information;
- involves a lot of data manipulation;
- systems can become very complex;
- assumes that lead times are constant and independent of the quantities ordered;
- ignores the order in which materials are really made and assumes that they are made in the order specified by the bill of materials;
- using MRP to schedule the production of parts can give poor results;
- the lot sizes suggested by MRP can be inefficient;
- MRP may not recognize capacity and other constraints;
- can be expensive and time-consuming to implement.

Despite these problems, MRP is widely used and standard software is readily available. You should, however, use this with caution. There are many organizations where MRP has failed to achieve expectations or to bring real benefits.

Perhaps the main reason for this is the mistaken belief that a system can simply be bought and switched on. MRP needs considerable changes to the way that an organization works, and these changes need commitment from all areas. Even if this commitment exists, it should not be assumed that MRP is simply an additional tool for planning. It is an integral part of the planning process and requires new procedures in many areas.

Overall, we can say that MRP may give considerable benefits, and many companies have made substantial savings, but to achieve these savings the companies have had to invest a lot of time and effort.

Summary

The main advantage of MRP is its ability to relate demand for materials directly to the master schedule. This can bring a number of benefits, ranging from reduced costs to better planning. The main disadvantages are the conditions that must be met, the complexity of systems and the inflexibility.

Review questions

9.6 What is the main advantage of MRP?

9.7 What is the main disadvantage of MRP?

9.8 What are the typical outputs from an MRP system?

Adjusting the MRP schedules

Batching orders

In the calculations that we have described, we placed an order exactly to cover the net requirement in a period. This would often suggest a series of frequent, small orders – perhaps with an order every period – which are inconvenient and expensive. It might be better to combine several of these small orders into a few larger batches. This is called *batching* or *lot sizing*, and there are four common approaches.

1. *Lot-for-lot* – where you order exactly the net requirement suggested by MRP for each period. This is the method we have used so far. It minimizes the amount of stock, but can give high ordering, set-up and administration costs.
2. *Fixed order quantity* – where you find an order size that is convenient, and always order this same amount. This might be a truckload, a container load, an economic order quantity, or some other convenient size.
3. *Periodic orders* – where you combine the net requirements over some fixed number of periods, and place regular orders for different quantities. An organization might, for example, place an order every month. Working to such a regular timetable gives the benefits of simplicity and making operations routine.

4. *Batching rules* – which use some specific procedure to calculate the best pattern of orders. Typically they look for the combination of orders that gives the minimum overall cost.

Because they aim for the lowest costs, batching rules are often the best approach. In practice, though, this can be a deceptively difficult scheduling problem, so organizations use simple rules that give reasonable results. We have already met one of these in Chapter 4, when we looked at discrete variable demand. Then we found the best policy of combining orders for the next N periods into one order, when N is the smallest number for which:

$$N \times (N + 1) \times D_{N+1} > \frac{2 \times RC}{HC}$$

where:

N = the period number in a cycle

D_{N+1} = demand in period $N + 1$ of a cycle

RC = reorder cost

HC = holding cost

To solve this, we use an iterative procedure, setting N equal to 1 and checking the values in the expression. If the inequality is invalid, it is cheaper to order for two periods than for one. This means that we have not found the best solution and set N equal to 2 and check the values in the expression. If the inequality is still invalid, the costs are still reducing and we keep on increasing N, until eventually the inequality will become valid. At this point we have found an optimal value for N.

Worked example

An MRP schedule gives the following pattern of demand for an item. If the reorder cost is £150 and the holding cost is £2 a week, find the costs of different batching policies over the 10 weeks.

Week	1	2	3	4	5	6	7	8	9	10
Demand	20	10	20	25	0	15	10	20	0	30

Solution

We can compare the costs of four different approaches to batching. The stock patterns for these are shown in Figure 9.10. The average stock in each period is the average of the opening and closing values. As you can see, the total costs are as follows.

	A	B	C	D	E	F	G	H	I	J	K
1	Batching rules										
2											
3	**Week**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
4	Demand	20	10	20	25	0	15	10	20	0	30
5											
6	**1. Lot-for-lot**										
7	Deliveries	20	10	20	25	0	15	10	20	0	30
8	Average stock	10	5	10	12.5	0	7.5	5	10	0	15
9	**Number of orders**		**8**		**Reorder cost**		**1200**				
10	**Stock holdings**		**75**		**Holding cost**		**150**		**Total cost**		**£1,350**
11											
12	**2. Fixed order quantity**										
13	Opening stock	0	28	18	46	21	21	6	44	24	24
14	Deliveries	48	0	48	0	0	0	48	0	0	48
15	Closing stock	28	18	46	21	21	6	44	24	24	42
16	Average stock	14	23	32	33.5	21	13.5	25	34	24	33
17	**Number of orders**		**4**		**Reorder cost**		**600**				
18	**Stock holdings**		**253**		**Holding cost**		**506**		**Total cost**		**£1,106**
19											
20	**3. Periodic**										
21	Opening stock	0	55	45	25	0	45	30	20	0	30
22	Deliveries	75				45				30	
23	Closing stock	55	45	25	0	45	30	20	0	30	0
24	Average stock	27.5	50	35	12.5	22.5	37.5	25	10	15	15
25	**Number of orders**		**3**		**Reorder cost**		**450**				
26	**Stock holdings**		**250**		**Holding cost**		**500**		**Total cost**		**£950**
27											
28	**4. Batching rule**										
29	Opening stock	0	30	20	0	25	25	10	0	30	30
30	Deliveries	50			50				50		
31	Closing stock	30	20	0	25	25	10	0	30	30	0
32	N*(N+1)*D	20	120	300	0	90	120	400	0	180	300
33	Average stock	15	25	10	12.5	25	17.5	5	15	30	15
34	**Number of orders**		**3**		**Reorder cost**		**450**				
35	**Stock holdings**		**170**		**Holding cost**		**340**		**Total cost**		**£790**

Figure 9.10 Comparison of batching rules for worked example

1. *Lot-for-lot*. Eight deliveries and an average stock level of 7.5 units give a total cost of $8 \times 150 + 75 \times 2 = £1{,}350$.
2. *Fixed order quantity*. The average demand is 15 units a week, so a useful approach is to calculate the economic order quantity as $\sqrt{(2 \times 150 \times 15/2)} = 47.4$, and round this up to 48. Then four deliveries and an average stock of 25.3 units give a total cost of $4 \times 150 + 253 \times 2 = £1{,}106$.
3. *Periodic*. If we assume a regular order once every 4 weeks, we get three deliveries and an average stock of 25 units, giving a total cost of $3 \times 150 + 250 \times 2 = £950$.
4. *Batching rule*. Using the rule we derived in Chapter 4 with $2 \times RC/HC = 2 \times 150/2 = 150$ gives three deliveries and a total cost of $3 \times 150 + 170 \times 2 = £790$.

In this case the batching rule works well and gives the lowest cost.

Closed-loop MRP

Combining several smaller orders for an item into fewer, larger ones can reduce costs, but they make the demand for following levels of materials even more sporadic. Then higher levels of materials have demands that become increasingly lumpy. This effect might be particularly severe when several products use the same materials, when MRP systems will link the different products and generate combined orders, but give very sporadic results. Operations and procurement have to deal with this and they can meet new problems. A supplier, for example, might have limited capacity and prefer level demand. They may not be able to cope with the peaks and troughs of sporadic demand unless they get adequate warning. It is obviously best to anticipate potential problems with bottlenecks during the planning stage, and then schedules and capacity can be adjusted before the plans are finalized. In other words, we introduce feedback between capacity planning and MRP. This is most useful when:

- proposed plans would break some capacity constraint; then the MRP system detects the overload and initiates early rescheduling. In this way, the MRP system takes an active part in capacity planning.
- operations are interrupted; the control system detects this and updates the master schedule. But it must work with the MRP system to see what production is feasible with the materials available. Now MRP forms part of the control system.

Systems with this kind of feedback are called *closed-loop MRP*. You can imagine a typical operation with feedback. The initial MRP run gives schedules for the supply of materials, and these are translated into required capacities. If available capacity cannot meet the demand, the system makes changes by adjusting capacity to meet the higher demand, rescheduling to make demand fit into the available capacity, or some compromise between these two. This third option is the most common with the system making a series of iterative adjustments, of the kind that we met with resource planning. These iterations are continued until a reasonable solution appears. If the adjustments to capacity and schedule cannot reach an acceptable solution, the master schedule may have to be redesigned. This approach of linking capacity planning to MRP is sometimes called *capacity requirements planning*, and is summarized in Figure 9.11.

Summary

Basic MRP schedules tend to give small, frequent orders. We can reduce costs by combining several of these into a single larger order using a batching rule. This can give increasingly sporadic demand that causes problems with capacity. These can be avoided by adding feedback in closed-loop MRP.

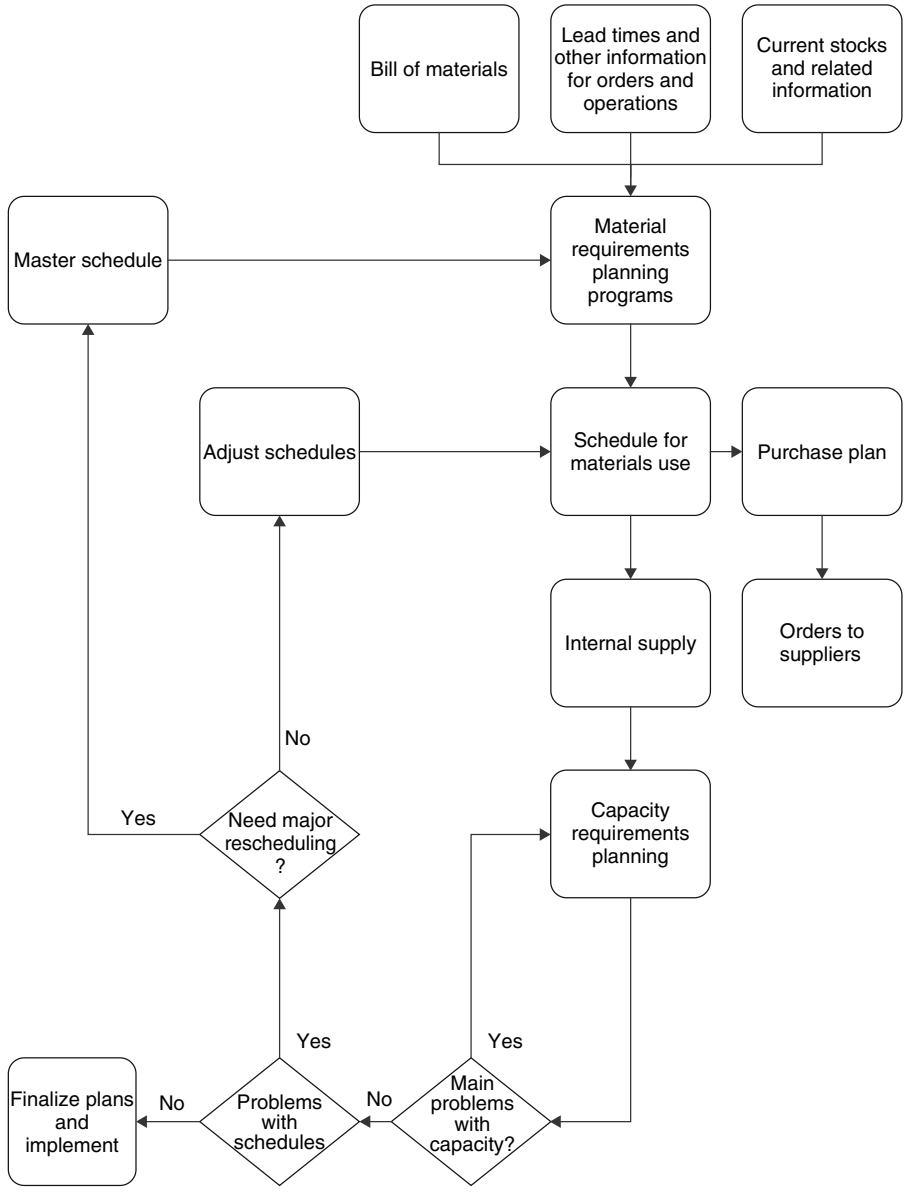


Figure 9.11 Closed loop MRP system

Review questions

9.9 Why do independent demand methods generally perform badly with intermittent demand?

9.10 It is always preferable to combine several, small orders into a few larger ones. Do you think this is true?

9.11 What is a batching rule?
9.12 What is the key element of closed-loop MRP?

Extensions to MRP

Initial ideas

When people started looking for improvements to the basic MRP approach they began by adding fairly small features to deal with variable supply, wastage, supplier reliability, poor quality, variable demand, variable lead times, and so on. Often these problems can be dealt with by simple rules. If, for example, an organization finds that the timing of deliveries is unreliable, it can add a safety time and order materials some time before they are actually needed; if quality is a problem, it can add some safety stock to replace materials that do not meet quality standards; if materials are cheap, it can use some minimum order size. Simple rules of this kind have been developed to deal with many different problems. They might raise costs, but this price is worth paying to guarantee uninterrupted operations – especially when you remember that inflexibility is a continuing problem with MRP.

Perhaps the next stage of extensions moved MRP away from manufacturers and towards other types of organization. It can be more difficult to apply in services, largely because they do not always have a detailed master schedules to define production. Nonetheless, many services have successfully adopted MRP. Hospitals, for example, use MRP to schedule surgical operations and make sure that supplies and equipment are ready when needed. They design a master schedule of planned surgical operations; the bill of material contains information about the facilities, equipment and resources needed for each type of surgery; and the inventory file contains information about theatres, staff, surgical instruments, disposable materials, reusable instruments, and so on. In the same way, hotels, universities, airlines, football clubs, conference organizers, film companies and a whole range of other services now use the principles of MRP.

After mentioning these initial adjustments, and the use of MRP by other types of organization, we come to the major extensions. We have seen how closed-loop MRP adds feedback to avoid potential problems with capacity but it really does much more than this. With capacity requirements planning it becomes a way of extending the MRP approach further into an organization. It extends MRP beyond the basic purchase of materials, and into more central areas of production and capacity planning. The next moves in this direction were again originated by manufacturers, and have been brought together under the umbrella of **manufacturing resource planning**.

Manufacturing resource planning

MRP gives schedules for the arrival of materials needed by operations, and closed-loop MRP makes sure that there is enough capacity actually to supply these materials. But materials are only one resource. As well as scheduling materials,

organizations have to schedule people, equipment, facilities, finances, logistics, and any other resource that they use. Surely we can use the same MRP approach for these other resources.

So we start with MRP schedules for the delivery of materials, and hence the timing of operations within the organization to supply the materials. But when we know the timetables for these operations, we can use them to schedule all the equipment needed. And when we have timetables for the equipment, we can use them to schedule the people who work on them. And when we have timetables for the people, we can use them to schedule other facilities, such as catering and transport. But we can go even further than this. If we know when we are getting deliveries, we can schedule transport; looking at the finished goods, we can plan delivery to customers and distribution. Continuing in this way, we could eventually build an integrated system that would 'explode' the master schedule to give timetables for all the jobs, equipment, operators, machines, facilities, etc. needed to achieve it. This integrated system for scheduling resources is *resource requirements planning* (actually, you have to be a bit careful here, because people sometimes use this phrase to mean different things).

In principle, there is no reason why we should limit the use of this MRP approach to scheduling basic resources. Why not, for example, continue and look at the associated finance, marketing, logistics, human resource management, etc.? Eventually we would have a completely integrated system that would use the master schedule as the basis for planning all the activities within an organization. This is the aim of *manufacturing resource planning* or *MRP II*.

- *MRP II* gives an integrated system for synchronizing all functions within an organization.
- It connects schedules for all activities back to the master schedule.

In principle, linking all activities to the master schedule can give very efficient operations. No unnecessary work is done, there are no delays because of late deliveries or shortages, no stocks of work in progress accumulate, and products move smoothly through the whole process. These benefits have encouraged many organizations to extend their MRP systems towards MRP II.

Although the idea of MRP II seems fairly straightforward, there are considerable practical difficulties. At a simple level, it is difficult to get schedules that everyone accepts as being good and workable. In common with most planning, MRP II does not stop at a single run to find the 'best' solution, but uses an iterative approach to find a reasonable solution that everyone accepts. It is clearly difficult to get agreement from so many people with different interests and aims, and the series of 'what-if' tests that look for improvements can become daunting. Jones and Towill (1998) describe these as 'monster systems'.

Perhaps the most serious problem with MRP II is the difficulty of getting complete integration of all functions and systems. Many organizations have asked if the effort needed to get such close integration is actually worth the rewards. We have already said that MRP systems tend to be inflexible and unable to respond

quickly to changes, so a whole organization run in this way might become cumbersome and unwieldy. Rather than give an efficient process, MRP II might leave it vulnerable to changing conditions. It is certainly fair to say that very few organizations work with full MRP II, and most are happier to implement parts of the system. They use different names for these partial systems, and a common is *distribution resource planning* or *logistics resource planning* where the MRP approach is used to plan logistics (again, unfortunately, people tend to use these terms rather inconsistently).

Distribution resource planning (DRP) plans the flow of products out to customers in the same way that MRP plans the flow of materials in from suppliers. Suppose, for example, that a local warehouse wants to make a delivery to a customer in four weeks' time, it knows when this delivery has to come from a regional warehouse. Then the regional warehouse knows when to arrange a delivery from a national warehouse, the national warehouse knows when to arrange a delivery from a supplier, and so on back through the supply chain. The problem comes with initiating this flow. With MRP this is done by the master schedule – as it should be with DRP. Then we effectively have customer demand used to design a master schedule, and this is used for both MRP and DRP. Sometimes, especially when there are long lead times, actual customer orders can be used to trigger DRP, giving a very close match between the supply and demand of logistics. Often, though, organizations cannot get rigid delivery plans so far in advance, and they have to base DRP on forecasts of likely deliveries. Provided these forecasts are accurate, DRP can give useful results, but they have lost some of the essential benefits of MRP.

Working with other organizations

The systems we have described so far are all used within an organization. But in Chapter 1 we discussed the benefits of integration along the supply chain. This allows organizations in a supply chain to improve their overall performance by co-operating, exchanging information and co-ordinating their operations. Now you can see the next extensions to MRP, which co-ordinate operations of different organizations within a supply chain. This is the basis of *Enterprise Resource Planning* (ERP).

ERP uses the characteristic approach of MRP, but now focuses on the whole supply chain. It extends MRP beyond a single organization, to include suppliers and customers (as shown in Figure 9.12). Suppose that a manufacturer ran its own MRP and found that its master schedule needed 100 units of some materials at the beginning of May. If its MRP system is linked to the supplier's system, the supplier can now start scheduling activities to make sure these materials are ready in time. Then the supplier, in turn, could link its system to its own suppliers, and so move backwards through the supply chain, creating an integrated planning system.

ERP obviously relies on complete trust between organizations, and a free flow of information. In principle, this is relatively easy to organize, using the Internet and e-commerce. However, you can imagine the complexity of systems needed, and the practical problems that arise. Although many systems suppliers will disagree, at the moment it is fair to say that ERP is still evolving.

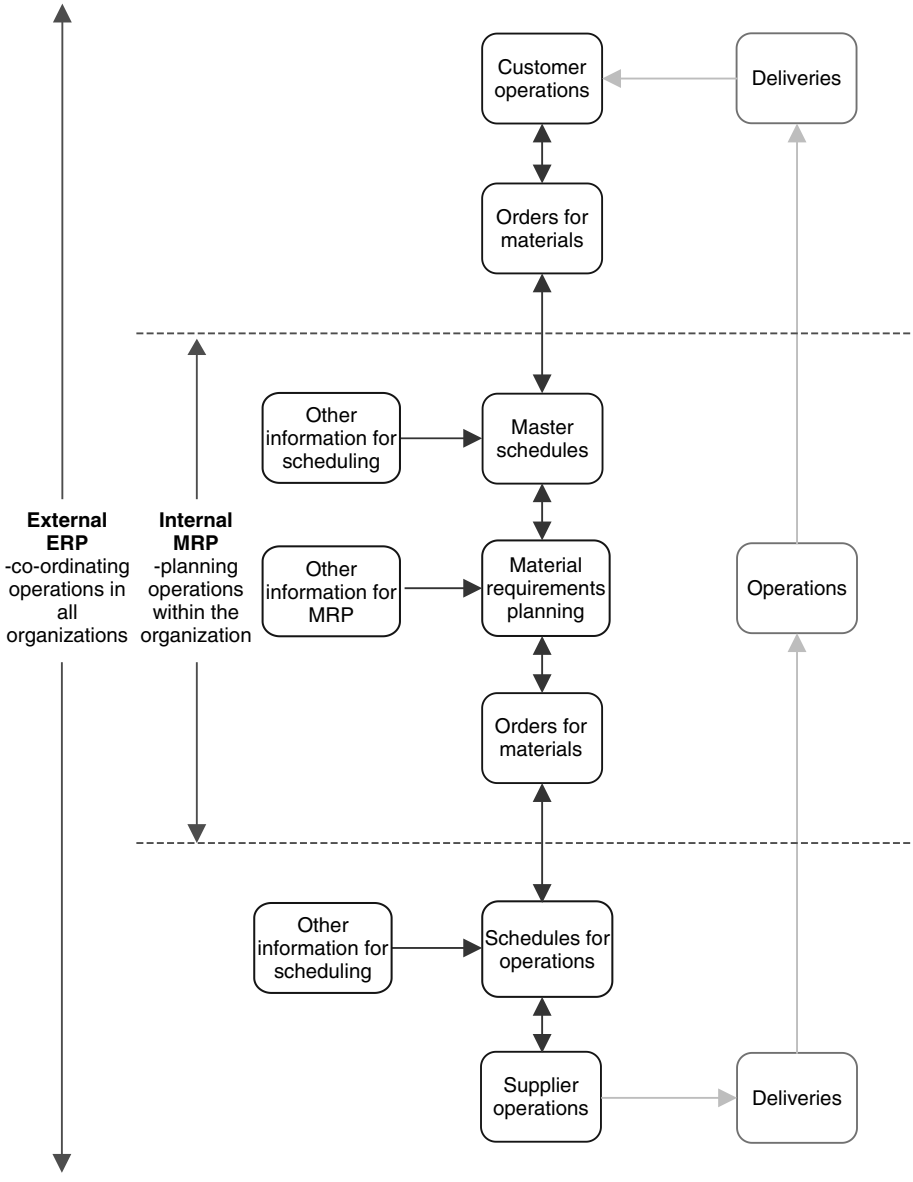


Figure 9.12 Enterprise resource planning

We can summarize the main stages in the development of the MRP approach as follows:

- Material requirements planning, where the schedules for materials are found from the master schedule.
- Adjustments to MRP to allow for batching, safety stocks, combined orders, etc.

- Closed-loop MRP (with capacity requirements planning) including feedback to make sure that the proposed schedule is feasible and presents no major problems (particularly with capacity).
- Resource requirements planning extends the MRP approach to other resources.
- Distribution (or logistics) resource planning schedules the flow of materials through a supply chain.
- Manufacturing resource planning extends the approach to include other functions within an organization.
- Enterprise resource planning extends the approach to other organizations in the supply chain.

Summary

There have been continual improvements to the basic MRP approach. Initial adjustments included batching of small orders, smoothing demand, adding safety stocks and iteratively improving schedules. More significant changes include feedback with closed-loop MRP, capacity requirements planning and DRP. MRP II extends the scope of the MRP approach to other functions, giving an integrated approach to planning. ERP extends the approach to other organizations in the supply chain.

Review questions

9.13 How could you extend the approach of MRP?

9.14 What exactly is MRP II?

9.15 Distribution resource planning relies on large and integrated computer systems. Do you think this is true?

Chapter review

- Despite their widespread use, independent demand methods have some weaknesses. We can modify the basic models, but in some circumstances – such as batch production – an alternative approach might be better. Material requirements planning gives a widely used alternative for dependent demand.
- Material requirements planning uses a bill of material to explode a master schedule and find the gross requirements for materials needed for production. Information about current stocks, orders outstanding, reserved stocks, etc. are then used to find order quantities. Lead times are use to time phase these orders so that materials arrive in time for use.
- The main advantage of MRP is its ability to relate demand for materials directly to the master schedule. This brings a number of related benefits, ranging from reduced costs to better planning.

- The main disadvantages of MRP are the conditions that must be met, the complexity of systems and the inflexibility.
- Basic MRP schedules tend to give small, frequent orders. We can reduce costs by combining several of these into a single larger order. We can also make other adjustments to smoothe demand, add safety stocks, iteratively improve schedules, and so on.
- MRP can give sporadic demand that causes problems with capacity. These can be avoided by adding feedback in closed-loop MRP. This leads to other extensions like capacity requirements planning and DRP.
- MRP II extends the MRP approach to other functions, giving an integrated approach to planning. ERP extends the approach to other organizations in the supply chain

Project

MRP needs specialized programs. What features would you expect to find in commercial MRP packages? Many companies provide software and consulting services to install MRP systems. Search the Web and see what suppliers you can find (you might start with SAP, a major supplier, on www.SAP.com). Do their products have the facilities you expected? Are there any apparent differences in their systems? Can you find evidence for real improvements brought by the systems?

Problems

9.1 Elsbeth Home Furnishing Ltd. assembles dining room tables using bought-in parts of four legs and a top. These have lead times of two and three weeks respectively, and assembly takes a week. The company receive orders for 40 tables to be delivered in week 5 of a production period and 80 tables in week 7. There are current stocks of four complete tables, 80 legs and 44 tops. When should the company order parts?

9.2 A master schedule shows 100 units of a product, AA, to be made in week 11 of a cycle, 220 units in week 14 and 150 units in week 16. There are currently 50 units of the product in stock. Each unit of the product takes a week to assemble from 3 units of part BB and 4 units of part CC. Each unit of part BB is made in one week from 2 units of material DD and 5 units of material EE. Part CC is assembled in 2 weeks from 3 units of component FF. Lead times for DD, EE and FF are 3, 2 and 1 weeks respectively. Current stocks are 200 units of BB, 250 of CC, 120 of DD, 400 of EE and 200 of FF. The company keeps minimum stocks of 20 units of AA, 40 units of DD, 100 of EE and 100 of FF. The minimum order size for EE is 500 units, while FF can only be ordered in discrete batches of 100 units. An order placed with a sub-contractor for 200

units of CC is expected to arrive in period 8. Design a timetable of activities for the company.

9.3 RTK makes three sizes of filing cabinet with 2, 3 and 4 drawers. Each cabinet consists of a formed case, drawers and a lock. Each case is made from drawer slides (two for each drawer) and a formed case (which is itself made from a sheet of steel). Each drawer is made from two roller supports, a handle and a formed drawer (which is made from a sheet of steel). The lead times, in weeks, are:

cabinet	1	lock	1	sheet steel	3
formed case	1	roller support	3		
drawer	2	handle	2	drawer slide	3

There are currently stocks of 30 complete drawers and 80 roller supports, and a delivery of 300 roller supports is expected in period 1. The master schedule for the next 10 weeks is as follows:

Week	1	2	3	4	5	6	7	8	9	10
2 drawer cabinets		50			100			100		100
3 drawer cabinets				60	120		80		120	
4 drawer cabinets			50		150		110			110

When would you schedule orders and production?

9.4 It costs €1.50 to store a unit of an item for one month. The total cost of placing an order for the item, including delivery, is €950. An MRP analysis has found the following demands for the item:

month	1	2	3	4	5	6	7	8	9	10	11	12	13
demand	50	25	30	30	50	60	40	30	20	30	40	50	70

Compare the costs of different ordering policies and say which is the best.

9.5 A company makes each unit of product APF83 from 12 units of BIL179, 10 units of CPX213 and 20 units of DDY303. Each unit of BIL179 is made from 2 units of EMA108, 2 units of FFT187 and 2 units of GHO951. Each unit of both EMA108 and DDY303 is made from 6 units of HOX229. A master schedule needs 120 units of APF83 to be ready by week 8 of a planning cycle and 100 units by week 10. There are minimum order sizes of 4,000 units for HOX229, and 1,000 units of both FFT187 and GHO951. Information about stocks and lead times in weeks (either for assembly or orders) is as follows.

Item	Current stocks	Minimum stocks	Lead time (weeks)
APF83	40	20	2
BIL179	460	100	3
CPX213	680	200	1
DDY303	820	200	3
EMA108	720	400	2
FFT187	1,240	400	2
GHO951	1,660	400	2
HOX229	2,400	400	4

Design a schedule for materials supply.

Discussion questions

9.1 Does MRP really give better results than independent demand approaches? Give some examples to support your views.

9.2 In what type of companies is MRP most widely used? Are there any specific problems with services? Give some real examples to support your views.

9.3 MRP links operations to the master schedule. But many operations are done more efficiently in a different order and with different batch sizes to those given by MRP. How do you think such problems can be overcome?

9.4 MRP is complicated and if you start linking other operations, systems, functions and organizations, the result becomes hugely complicated. The ideas of MRP II and ERP seem good in principle, but the systems become so unwieldy that they can never work properly. Even if they did work, operations would be too inflexible and burdensome to cope with more agile competitors. Do you think that this is true?

9.5 Distribution resources planning is based on the mistaken view that deliveries to customers can be planned with the same certainty as operations in production. Can DRP ever work as well as MRP?

9.6 MRP has evolved from a simple method of calculating order times, to an integrated system for controlling all functions along the supply chain. What are the most important steps in this evolution? What are the next steps?

References and further reading

Baudin, M. (1990) *Manufacturing Systems Analysis*, Englewood Cliffs, NJ: Prentice-Hall.

Browne, J., Harhen, J. and Shivnan, J. (1996) *Production Management Systems*, Reading, MA: Addison-Wesley.

Drexl, A. and Kimins, A. (eds) (1998) *Beyond Manufacturing Resource Planning*, Berlin: Springer-Verlag.
Jones, D. and Towill, D. (1998) New directions in logistics research, *Control*, March, pp. 15–19.
Luber, A.D. (1995) *Solving Business Problems with MRP II*, London: Butterworth-Heinemann.
Orlicky, J. (1975) *Material Requirements Planning*, New York: McGraw-Hill.
Sheikh, K. (2002) *Manufacturing Resource Planning*, New York: McGraw-Hill.
Ptak, C.A. (1997) *MRP and Beyond*, New York: McGraw-Hill.
Vollman, T.E., Berry, W.L. and Whybark, D.C. (1996) *Manufacturing Planning and Control Systems* (4th edition), New York: McGraw-Hill.
Wight, O.W. (1974) *Production and Inventory Management in the Computer Age*, Boston: Cahners Publishing.
Wight, O.W. (1982) *The Executive's Guide to Successful MRP II*, Williston, VT: Oliver Wight Publications.
Wight, O.W. (1984) *Manufacturing Resource Planning: MRP II*, Essex Junction, VT: Oliver Wight Publications.

10

Just-in-Time

Aims of the chapter

The last chapter described material requirements planning, which is a dependent demand method that finds demand for materials directly from the master schedule. This chapter continues the theme of dependent demand by introducing just-in-time. The distinctive feature of just-in-time is that it eliminates waste by organizing operations to occur at exactly the time they are needed. In this sense, stock becomes a waste of resources that should be eliminated.

After reading this chapter you should be able to do the following:

- appreciate the view that stock is a waste of resources that should be eliminated;
- describe the aims of just-in-time operations and the effects these bring to an organization;
- use *kanbans* to control JIT operations;
- discuss some important requirements of JIT, including quality management and relations with suppliers and employees;
- list the benefits and disadvantages of JIT;
- see how to extend JIT to other organizations in the supply chain;
- compare JIT with other methods of inventory control.

This chapter will emphasize:

- *just-in-time*, when operations are done just as they are needed;
- *pull systems*, which pull materials through a process;
- *kanbans*, for controlling JIT operations.

Principles of just-in-time

Just-in-time or *JIT* organizes all operations so they occur at exactly the time they are needed. They are not done too early (which would leave products and materials

hanging around until they were actually needed) and they are not done too late (which would give poor customer service). You can see this effect when you order a taxi to collect you at 12:00. If the taxi arrives at 11:30 you are not ready and it wastes time sitting and waiting; if it arrives at 12:30 you are not pleased and will not use the service again. When the taxi arrives at 12:00 – just-in-time for your trip – it does not waste time waiting, and you are pleased that the service arrives exactly when you wanted it.

With *just-in-time* operations, activities occur at exactly the time they are needed.

JIT and stocks

Although it seems a simple idea, JIT needs considerable changes to the way that organizations work and to management thinking. This effort can be rewarded by dramatic improvements in performance. Xerox was one of the early users of JIT and Jacobson and Hillkirk (1986) describe how 'Just-in-time purchasing at Xerox reduced the number of suppliers from 5000 to 400, cut stocks of copier parts by $240 million, and automated its warehouses'.

JIT is much more than a way of timetabling operations, and its supporters describe it as 'a way of eliminating waste' or 'a way of enforced problem solving'. We will illustrate this by its effects on stocks. Here, JIT immediately sees stock as a waste of resources that serves no useful purpose, with materials just sitting and waiting to be used. So it looks for ways of removing this waste. The traditional view has stock providing an essential buffer between operations, and managers asking how to provide this buffer at minimum cost. JIT changes the main question of inventory control to, 'How can we eliminate the need for stocks?'

JIT says that the answer to this comes by solving the problems that make stock necessary. If stocks are held to cope with variations in supply, the answer is to find ways of reducing the variation. If stocks are held to cover uncertain demand, the answer is to remove the uncertainty. So JIT identifies the problems that arise if stocks are eliminated, and then looks for ways of solving these problems.

In the last chapter we saw how MRP takes a step in this direction. It reduces the uncertainty in demand for materials by linking this to planned production. The result is lower average stock levels. In practice, MRP can only be used in certain circumstances and it still keeps some stocks to allow for problems with operations or deliveries. However, the more closely we can match supply to demand, the less stock we need to cover any differences. If we can completely eliminate any mismatch, we do not need any stocks at all (as illustrated in Figure 10.1).

By co-ordinating supply and demand so that they are closely matched, just-in-time can eliminate the need for stocks.

You can see an example of just-in-time operations when you buy fuel for a lawnmower. If your lawnmower has a petrol engine, there is a mismatch between the fuel supply that you buy from a garage, and demand when you actually mow

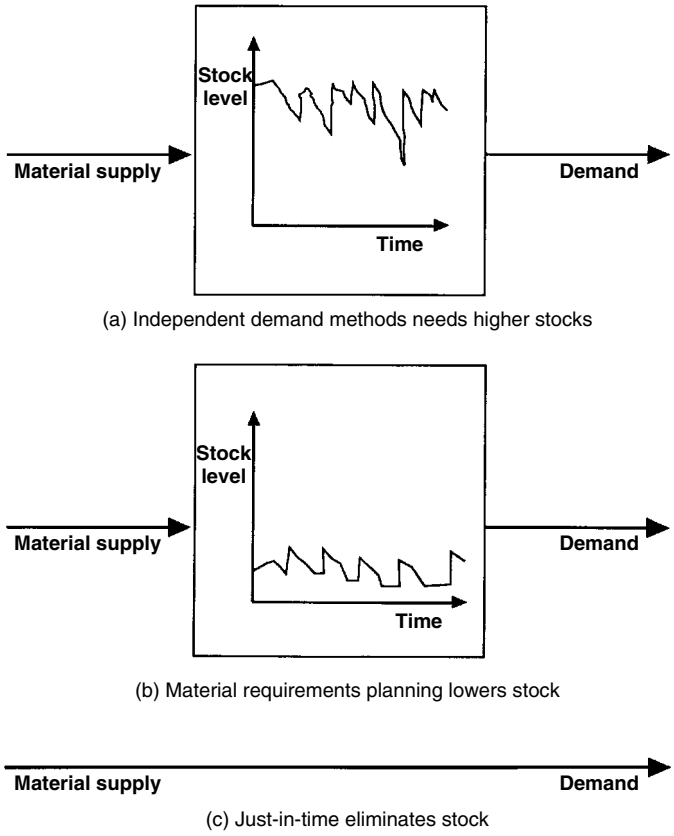


Figure 10.1 Comparison of stock levels with different approaches to inventory control

the lawn. You allow for this mismatch by keeping stocks of fuel in the petrol tank and spare can. This is the traditional approach to inventory control, where stocks are high enough to cover any likely demand. If your lawnmower has an electric motor, the supply of electricity exactly matches demand and you do not need any stocks. This is a just-in-time system.

Now you can imagine JIT in practice by thinking of a car assembly line. Just as the chassis moves down the line to one stage, an engine arrives at the same point and is fitted. This is repeated for all parts. As the car body arrives at another stage, four doors also arrive and are fitted. All the way down the line materials arrive just at the time they are needed, and the car is assembled in one smooth process.

So what happens when there really is a mismatch between supply and demand? What does a supermarket do when it sells loaves of bread one at a time, but gets them delivered by the truckload? The traditional answer is to hold enough stock to cover the mismatch – effectively the supermarket must store bread until it is sold or goes stale. JIT says that this is a mistake, and the alternative is to remove the mismatch. The supermarket could solve its problem by having smaller, more frequent deliveries, or opening a small bakery on the premises.

As you can see, JIT is based on simple principles. Instead of holding stocks to allow for problems, you look at the problems and solve them. We can summarize its argument about stock as follows:

- Stocks are held in an organization to cover short-term variation and uncertainty in supply and demand.
- These stocks serve no useful purpose – they only exist because poor co-ordination does not match the supply of materials to the demand.
- As long as stocks are held, there are no obvious problems and no incentive for managers to improve the flow of materials.
- Then operations continue to be poorly managed, with problems hidden by stocks.
- The real answer is to improve operations, find the reasons for differences between supply and demand, and then take whatever action is needed to overcome the differences.

Japanese manufacturers effectively developed this approach, and companies like Toyota spent years perfecting the methods (Shingo, 1981). Their work started in the 1970s and by the early 1990s half of European manufacturers used some form of JIT. Now it is fair to say that all major organizations use some elements of JIT. People have given the basic approach a variety of names, including stockless production, zero inventory, lean production, Toyota system, Japanese manufacturing, world-class manufacturing and continuous flow manufacturing.

The ideas behind JIT are clearly not new. In the 1920s iron ore arriving at Ford's plant in Detroit was turned into steel within a day and into finished cars shortly afterwards; McDonald's have made billions of hamburgers using elements of just-in-time; television news services collect reports just-in-time for their transmissions. The difference is that Toyota, and others, showed how JIT principles can be extended to all operations. They showed practical ways of running whole organizations based on JIT principles.

Summary

Just-in-time organizes operations to occur at exactly the time they are needed. This approach is based on an aim of eliminating all waste. Stocks are seen as a waste of resources that can be eliminated by co-ordinating supply and demand. The characteristic approach is to identify problems and then solve them, rather than hiding them under excessive stock.

Review questions

10.1 What is the main feature of JIT?

10.2 How does JIT's approach to inventory management differ from other methods?

10.3 The main benefit of JIT is that it reduces stocks of work in progress. Do you think this is true?

Main features for stocks

Changes of view

JIT is not just a way of minimizing stocks, but is a whole way of viewing operations. In this wider sense, JIT sees an organization as having a series of problems that hinder efficient operations. These problems include long lead times, unbalanced operations, constrained capacity, equipment breakdowns, defective materials, interrupted operations, unreliable suppliers, poor quality, too much paperwork and too many changes. Managers try to overcome the effects of these problems by holding large stocks, buying extra capacity, keeping back-up equipment, employing 'trouble-shooters', and so on. But these methods do not solve the problems, they only cover them up. A much more constructive approach is to identify the real problems and solve them. This approach leads to a number of changes in viewpoint:

- *Stocks*. As we have seen, organizations hold stocks to cover short-term differences between supply and demand. JIT assumes these stocks are actually used to hide problems. Organizations should find the reasons for differences between supply and demand, and then take whatever action is needed to remove them.
- *Quality*. Historically, organizations have defined some arbitrary level of acceptable quality for their products, such as, 'two defective units in 100 means the quality is acceptable'. JIT recognizes that all defects have costs and prevent smooth operations, so it is better to make sure that no defects are produced. This is the principle behind total quality management (TQM).
- *Suppliers*. We have stressed the need for co-operation between suppliers and customers, but this is particularly important with JIT. Organizations using JIT rely totally on their suppliers, so they cannot allow any kind of friction. Instead they emphasize their common objectives and work closely together, preferably in long-term partnerships and alliances.
- *Batch size*. Operations are often more efficient with large batch sizes, as they reduce set-up costs and disruptions. But if demand is smaller, the big batches give stocks that are held for a long time. JIT looks for ways of reducing the batch size so that it more closely matches demand.
- *Lead times*. Lead times are often fixed by suppliers, and they can be unnecessarily long. This reduces flexibility and encourages high stocks to cover uncertainty before another order can arrive. JIT looks for ways of continually reducing lead times.
- *Reliability*. When equipment breaks down, most organizations transfer resources and start making another product. JIT is based on continuous,

uninterrupted production, so it does not allow this kind of flexibility. Managers are forced to recognize the problem with reliability, they find the reason for the breakdown, and make sure it does not happen again.

- *Employees*. Some organizations still have a friction between 'managers' and 'workers'. JIT argues that this is a meaningless distinction. The welfare of everyone depends on the success of the organization, so all employees should be treated fairly and equitably.

By now, you can see that JIT is not just as a way of minimizing stocks, but is a whole way of viewing operations. One other key element of JIT is its view of administration as an overhead that is largely wasted. So it tries to simplify operations and systems so that the effort needed to control them is minimized. JIT systems are largely manual with little paperwork or decisions made away from the shop floor, which is in marked contrast to MRP, which is computerized, expensive to control, and based on decisions made by planners who are some distance from the shop floor.

This aim of simplicity means that the methods used by JIT are all practical and based largely on common sense. Thus plant layouts are simplified, routine maintenance is scheduled to avoid break-downs, everyone is trained in quality control to reduce the number of defects, simpler designs are used to reduce processing time, equipment set-up is changed to make it faster, reorder costs are reduced to allow smaller deliveries, suppliers are encouraged to make more frequent deliveries, and so on. These changes may seem fairly obvious, but they have major effects on operations. It is easy to say, for example, 'Suppliers are encouraged to make more frequent deliveries', but this can be very difficult to organize in practice. Such changes cannot be introduced in one go, but evolve with small continuous improvements over a long period. Toyota is reputed to have made continuous improvements in its operations for over 25 years before JIT was working properly (see Womack *et al*., 1990; Monden, 1994).

Types of operations

Unfortunately, JIT only works well in certain types of organization. The most successful users are probably car assembly plants that make large numbers of similar products in a continuous process. You can see why this is, from the following arguments.

- Every time there are changes to a process, or it switches from making one product to making another, there are delays, disruptions and costs. JIT says that these changes waste resources and should be eliminated. In other words, JIT needs a stable environment where a process makes large numbers of a standard product, at a fixed rate, for a long time.
- This stable environment can reduce costs by using specialized automation. Then JIT works best with high volume, mass production.

- The level of production must allow a smooth and continuous flow of products through the process. Each part of the process should be fully utilized, so the process is likely to be a well-balanced assembly line.
- Deliveries of materials are made directly to the assembly line at just the time they are needed. Suppliers must be able to adapt to this kind of operation. It would be impractical to bring each individual unit from suppliers, so the next best thing is to use very small batches.
- If small batches are used, reorder costs must be reduced as much as possible or the frequent deliveries will be too expensive.
- Lead times must be short or the delay in answering a request for materials becomes too long. This means working closely with suppliers, and even having them build facilities that are physically close.
- As there are no stocks to give safety cover, any defects in materials would disrupt production. Suppliers must, therefore, be totally reliable and provide materials that are free from defects.
- If something goes wrong, people working on the process must be able to find the cause, take the action needed to correct the fault, and make sure that it does not happen again. This needs a skilled and flexible workforce that is committed to the success of the organization.

We can continue arguing in this way and arrive at a list of the key elements in JIT operations. These include:

- a stable environment
- standard products with few variations
- continuous production at fixed levels
- automated, high volume operations
- a balanced process that uses resources fully
- reliable production equipment
- minimum stocks
- small batches of materials
- short lead times for materials
- low set-up and delivery costs
- efficient materials handling
- reliable suppliers
- consistently high quality of materials
- flexible workforce

- fair treatment and rewards for employees
- ability to solve any problems
- an efficient method of control.

Summary

JIT needs many changes to the way that organizations view their operations. Even when they are willing to make these changes, JIT does not work well in all circumstances. It is generally most successful with large-scale assembly.

Review questions

10.4 What type of process is JIT most suited to?

10.5 JIT cannot be used for small service operations. Do you think this is true?

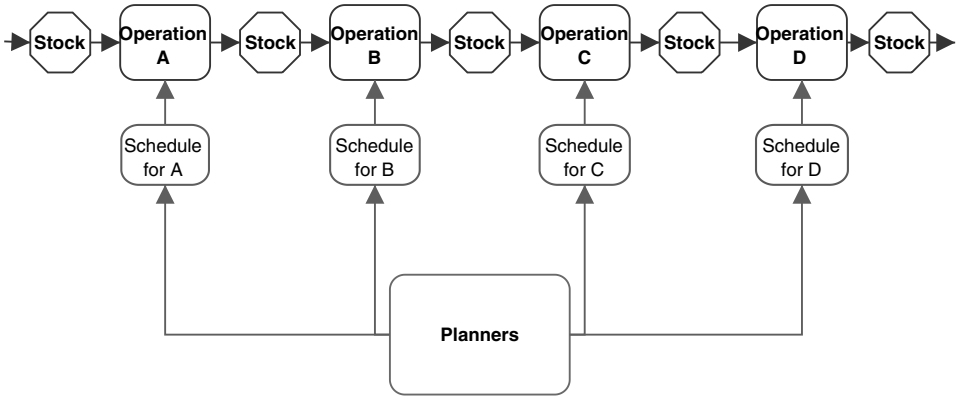
Achieving just-in-time operations

Push and pull systems

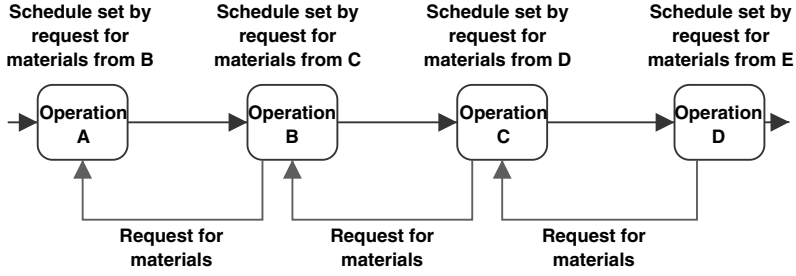
The key point of JIT is not just that it specifies an aim of organizing operations to occur exactly when they are needed, but also shows how this aim can be achieved. It develops a distinctive system that 'pulls' materials through the process.

In a traditional process, each operation has a timetable of work that must be finished in a given time. Finished items are then 'pushed' through to form a stock of work in progress in front of the next operation. You can imagine this with some carpenters whose schedule tells them to finish making a set of chairs by 15:00 next Thursday and then to pass them on to the painters. Unfortunately, this ignores what the next operation is actually doing – it might be busy working on something completely different, or be waiting for a different item to arrive, or be stopped by equipment failure. At best, the second operation must finish its current job before it can start working on the new material just passed to it. The result is delays and increased stock of work in progress.

JIT uses another approach to 'pull' work through the process. When one operation finishes work on a unit, it passes a message back to the preceding operation to say that it needs another unit to work on. The preceding operation only passes materials forward when it gets this request. This kind of process does not have earlier operations *pushing* work through, but has later operation *pulling* it through (as shown in Figure 10.2). You can see the difference in a take-away sandwich bar. With the traditional push system, someone makes a batch of sandwiches and delivers them to the counter where they sit until someone buys them. With a JIT pull system, a customer asks for a particular type of sandwich,



(a) **Stocks of materials with traditional 'push' planning**


Schedule set by request for materials from B
Schedule set by request for materials from C
Schedule set by request for materials from D
Schedule set by request for materials from E
Operation A
Operation B
Operation C
Operation D
Request for materials
Request for materials
Request for materials

(b) **Material 'pulled' through with JIT**

Figure 10.2 Comparison of 'push' and 'pull' planning

and this is specially made and delivered. This eliminates the stocks of work in progress.

In practice, there is inevitably some lead time between a later operation requesting material and having it arrive. So JIT arranges for messages to be passed backwards this lead time before they are actually needed. Materials are also delivered in small batches rather than continuous amounts. This means that JIT still has some stocks of work in progress, but these are much smaller than for equivalent 'push' systems. It is, though, fairer to say that JIT minimizes stocks rather than eliminates them.

Kanbans

Having defined a 'pull' system, JIT needs some way of controlling the flow of materials through a process. This is usually done by *kanbans* ('*kanban*' is the Japanese for a card, or some form of visible record). There are several ways of using *kanbans*, but they are based on the same principles.

- A message is passed backwards, asking preceding operations to send materials.
- Materials are only moved in standard containers which hold specific amounts.

- Only one container-full is produced or moved at a time.
- The size of a container is typically around 10 per cent of daily needs, but is the smallest reasonable batch that can be made.
- Containers can only be moved if the container has a *kanban* attached.
- A specific number of containers and *kanbans* is used, and this gives a rigid way of controlling the movement of materials.
- While it is simple to administer, this system makes sure that stocks cannot accumulate.

The simplest way of using *kanbans* is shown in Figure 10.3. Here, operation B uses the material from a container. When it needs more materials (i.e. when its stock of materials falls to some predetermined reorder level) a *kanban* is attached to an empty container and this is taken to the preceding operation A. Here the *kanban* is attached to a full container, which returns to operation B. Operation A now has an empty container, which is its signal to start work on this material, and it produces just enough to refill the container.

One obvious problem is that operations must be perfectly balanced, with the output from each operation exactly matching the requirements of following stations. If there is any imbalance some equipment remains idle until it is called on to start production. In practice, this problem is met in all operations and is by no means unique to JIT systems. However, with its emphasis on solving problems, JIT considers any imbalance to be unacceptable and looks for ways of eliminating it.

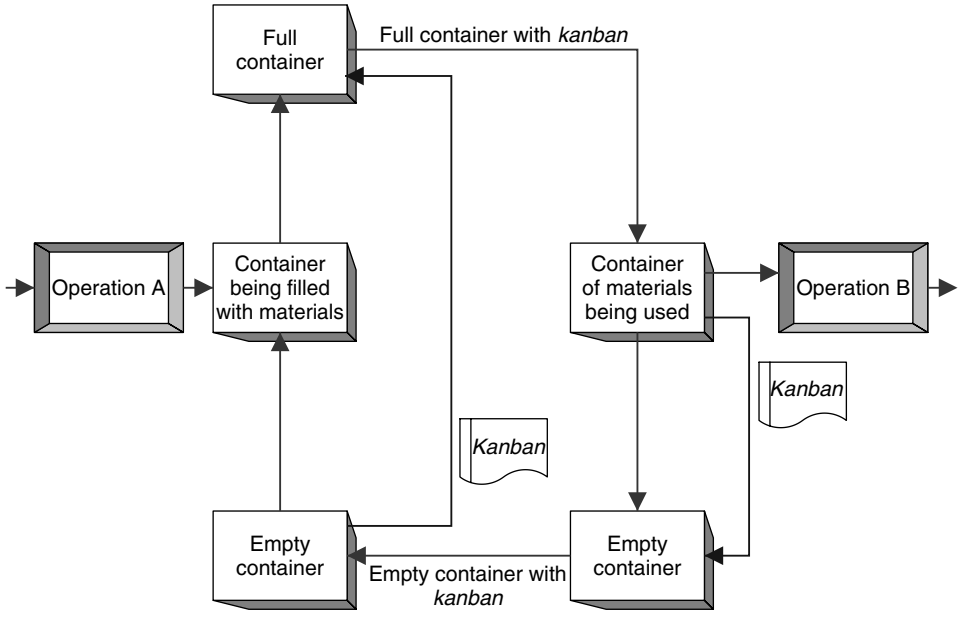


Figure 10.3 A simple *kanban* system

VELVSHER

The single card *kanban* system described is easy to implement, but it is rather limited and the most common system (shown in Figure 10.4) is slightly more complicated. This uses two distinct types of card, a production *kanban* and a movement *kanban*.

- When one operation needs more materials, a movement *kanban* is put on an empty container. This gives permission to take the container to a small stock of work in progress.

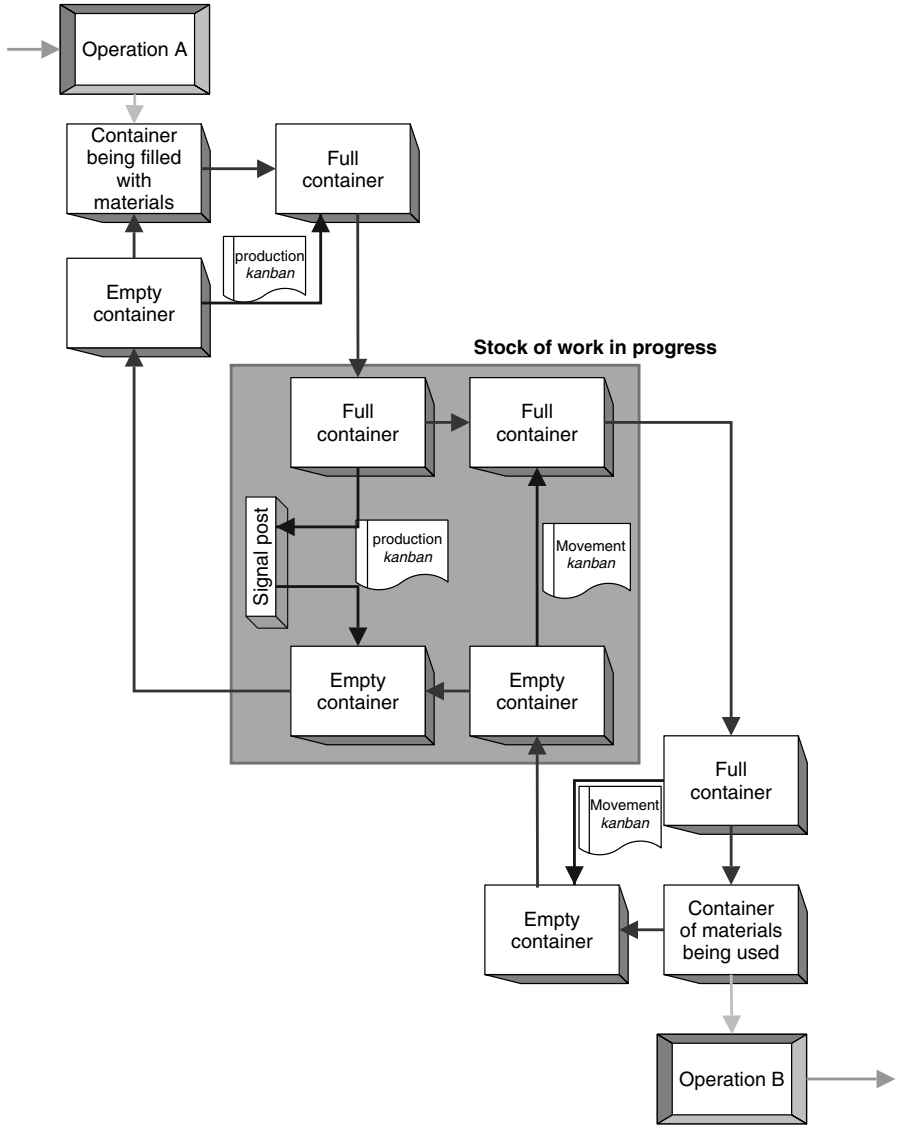


Figure 10.4 A common two card *kanban* system

- A full container is then found, which has a production *kanban* attached.
- The production *kanban* is removed and put on a post. This gives a signal for the preceding operation to make enough to replace the container of materials.
- A movement *kanban* is put on the full container, giving permission to take it back to the operation.

Although this system has a stock of work in progress, this stock is small. When a full container is removed, it is usually the only container in stock – and the materials are not replaced until the previous workstation makes them. JIT almost always uses a product layout – such as an assembly line – so this stock of work in progress is really a small amount that is kept in the line, and there is no actual movement.

Each full container in the store has a production *kanban* attached to it, so the number of *kanbans* effectively fixes the amount of work in progress. If there is only one production *kanban*, it means that the stock of work in progress is limited to at most one container of items. If there are two production *kanbans*, this doubles the stock of work in progress, and more *kanbans* would give even higher stocks. The aim of JIT is to work with minimum stocks and, therefore, a minimum feasible number of *kanbans*.

There are many variations on these themes. Some systems use different *kanbans* for emergency requests, high priority needs, materials requested from suppliers, signals for batch processes to start, and so on. The amount of information given on each *kanban* can vary from almost nothing to detailed instructions for production. Many organizations use electronic *kanbans*, coloured balls, bar codes and optical character readers. Bicheno (1999) says that, 'There are many types of kanbans . . . ''kanbans for all seasons'' '. Whatever the differences in detail, each system is based on a signal passed backwards from one stage in a process to the previous stage, to show when it is time to start operations.

Setting the number of *kanbans*

Suppose demand for an item is D per unit of time. Each container holds C units of an item (and as a guideline the value of C should be below 10 per cent of expected daily demand). Each container spends:

- a time TP in a production part of a cycle (waiting, being filled and moving to the store of work in progress);
- a time TD in a demand part of a cycle (waiting, being emptied and moving to the store of work in progress).

Then the total cycle length is TP + TD and the number of *kanbans* to maintain smooth operations is:

$$\text{Number of } \textit{kanbans} = \frac{\text{demand in the cycle}}{\text{size of each container}}$$

or

$$K = \frac{D \times (TP + TD)}{C}$$

Any number of *kanbans* above this adds an element of safety. When a new JIT system is installed, the organization might keep some flexibility by having a fairly large number of *kanbans*. Over time, though, JIT will look for continuous improvements and will reduce the number to a minimum. As a guide, the maximum element of safety should be less than 10 per cent so that:

$$K < \frac{D \times (TP + TD) \times (1 + SF)}{C}$$

where SF is a safety factor of less than 0.1. Then the maximum stock of work in progress is:

$$\text{Maximum stock level} = K \times C = D \times (TP + TD) \times (1 + SF)$$

Worked example

Demand for an item is 100 units an hour. The item is moved in containers that hold 10 units. Each container spends an average of 15 minutes in the production part of a cycle, and 30 minutes in the demand part. What is the minimum number of *kanbans* needed to control stocks? What is the resulting stock of work in progress? Improved operations can reduce the time in the demand part of the cycle to 20 minutes, but initially this needs a safety factor of 10 per cent. How many *kanbans* does this need?

Solution

The figures given, in consistent units, are:

$$D = 100 \text{ units an hour}$$
$$TP = 0.25 \text{ hours}$$
$$TD = 0.5 \text{ hours}$$
$$C = 10 \text{ units}$$

With no margin for safety the number of *kanbans* needed is:

$$K = \frac{D \times (TP + TD)}{10} = \frac{100 \times (0.25 + 0.5)}{10} = 7.5$$

K must be rounded up to 8, giving a stock of work in progress of:

$$K \times C = 8 \times 10 = 80 \text{ units}$$

Reducing TD to 20 minutes, and adding a 10 per cent safety margin gives:

$$K = \frac{D \times (TP + TD) \times (1 + SF)}{C} = \frac{100 \times (0.25 + 0.33) \times 1.1}{10} = 6.4$$

K must be rounded up to 7, giving a stock of work in progress of:

$$K \times C = 7 \times 10 = 70 \text{ units.}$$

Summary

JIT needs a simple means of controlling the supply of materials. *Kanban* is the usual method, using cards to signal the pull of demand from later operations to earlier ones. There are many different ways of using *kanbans*.

Review questions

10.6 What is the difference between 'push' and 'pull' systems?

10.7 What is the purpose of a *kanban?*

10.8 How is the amount of work in progress limited in JIT?

Other effects of JIT

We have already mentioned some of the features that you would expect to see in an organization using JIT. It has an effect on everything, from the way that goods are ordered to the quality of products. It is, therefore, a step that needs total commitment from everyone in the organization. Because they are so important for inventory management, we will mention four of the key elements of JIT again. These are the relations with suppliers, quality management, authority to stop operations, and respect for employees.

Relations with suppliers

Traditionally there has been some friction between organizations and their suppliers. Organizations pay money to suppliers, and many people think that one can only benefit at the expense of the other. If an organization gets a good deal, it automatically means that the supplier is losing out – if the supplier makes a good profit, it means that the organization pays too much. We said in Chapter 1 that this friction causes serious problems with loyalty and co-operation. Suppliers set rigid conditions, and as they have no guarantee of repeat business, they try to make as much profit from each sale as possible; organizations shop around to make sure they get the best deal, and remind suppliers of the competition. The result is uncertainty about orders, constant changes in an organization's suppliers

and customers, changing products, varying order sizes, different times between orders, uncertainty about repeat orders, changes in the costs, and so on.

Organizations using JIT rely on their suppliers being completely dependable and removing all uncertainty from supply. The only way of achieving this is through co-operation, recognizing that organizations and their suppliers both want a mutually beneficial trading arrangement. If they can agree conditions that satisfy both the organization and the supplier – with both feeling that they get the best possible deal – this is far better than working with unnecessary friction. The implication is that organizations should identify the best suppliers and always order from them. This reinforces the ideas behind strategic alliances and partnerships that we met in Chapter 1.

Without close co-operation it can be difficult for suppliers to make the small, frequent deliveries needed, and to co-ordinate their deliveries with demands. Sometimes they respond by increasing their stock of finished goods to ensure the required pattern of delivery. The effect of this is simply to move stocks from an organization's raw materials store to the suppliers' finished goods store. This does not reduce overall costs and might even increase them. The aim of JIT is to eliminate stocks rather than move them to another point in the supply chain. And, again, the way to achieve this is through co-operation.

JIT recognizes the benefits of stability. Any changes to products or operations inevitably cause difficulties with suppliers, and these problems expand as they pass back through the supply chain. A small change in a finished product has major effects on earlier suppliers. These difficulties are inefficient, so JIT relies on a product that remains largely unchanged through long production runs. This stability allows suppliers to specialize in one type of item, and they may reduce their product range and number of customers. Many suppliers become 'focused factories', which are small plants that concentrate almost entirely on making one product but aim to make this very well and very efficiently. They focus all their operations on making a single product very efficiently, and their small size makes them easy to manage and less expensive to run.

They can also share information without the threat that this will be used to get some form of trading advantage, and they can share ideas about products and designs. Suppliers can make suggestions for improvements, without the fear that their future profits will be reduced. Long-term agreements also reduce the threat that a dominant organization will look for vertical integration and either buy out a supplier or start making materials itself.

Ideally, suppliers will work so closely with an organization that they install JIT in their own operations and become part of an extended JIT system (as suggested in Figure 10.5). Whenever an organization needs some material it sends a vehicle with empty containers and *kanbans* to the supplier. The empty containers are exchanged for full ones from the supplier's stock of finished goods; the *kanbans* are transferred to the full containers, which are returned to the customer. The supplier now has empty containers, which show that it is time to replace the contents.

Of course, the close relationships between an organization and its suppliers can also bring problems. JIT is very demanding of suppliers, insisting on perfect quality, small orders, frequent deliveries, short lead times, shared information,

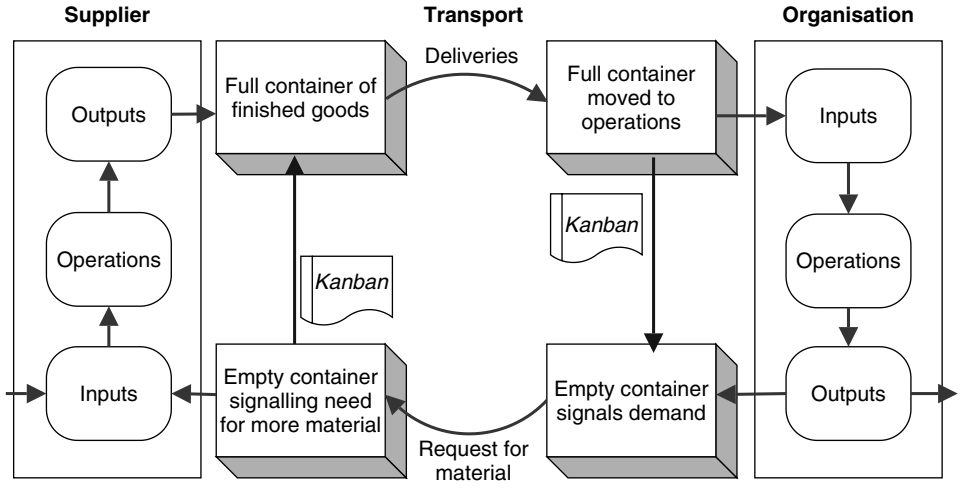


Figure 10.5 Integrating a supplier into the JIT system

simple control systems, little paperwork, nearby locations, low costs, continuing improvement to products and operations, ideas for better products, extension of JIT to suppliers, and so on. These conditions would be difficult even for a supplier with only one customer, but most suppliers have to consider many different customers with different needs. In reality, there has to be some compromise on both sides to give continuing benefits.

Jidoka – quality at source

JIT does not allow safety stocks, so materials delivered must be perfect and free from any defect. There are two ways of achieving this. First, the organization can accept the possibility of defects and check the quality of every unit as it arrives. But this is expensive and time-consuming, so it destroys many of the benefits that JIT is aiming for. The second alternative is to ensure that all items arriving are of perfect quality. This is clearly the better alternative and is the one adopted by JIT.

Traditionally manufacturers have accepted a certain proportion of deliveries as defective. Typically, companies would quote a maximum defective rate in their products, typically around 1 per cent, and they would sometimes compensate for this by delivering an extra amount of 'overage'. One serious problem, though, is that even low rates of defects are not good enough for JIT. If 99.9 per cent of materials have acceptable quality, the probability that a unit made from 100 parts is acceptable is 0.999^{100}, which is about 0.9 – meaning that 10 per cent of units are defective.

JIT needs perfect quality, and the way to achieve this came from Japan. The methods have become known as *total quality management* or *TQM*, and in recent years they have had such an effect that some people refer to a 'quality revolution'. This happen for four main reasons:

1. Improved processes can make products with guaranteed high quality.
2. High quality gives producers a competitive advantage.
3. Consumers have become used to high quality products, and will not accept anything less.
4. High quality reduces costs.

Perhaps the most interesting comment here is that high quality gives lower costs. At first this seems to go against the traditional view that higher quality can only be bought at higher cost, but if you look at the wider costs, you see that some actually go down with increasing quality. Imagine that you buy a washing machine that is faulty. You complain, and the manufacturer arranges for the machine to be repaired. The manufacturer could have saved money by finding the fault before the machine left the factory and it could have saved even more by making a machine that did not have a fault in the first place. If we look at the costs of quality more closely, we can divide them into four types.

1. *Prevention costs* are incurred to prevent defects happening. They include direct costs for the product itself, such the use of better materials, designing for ease of production, inclusion of features to ensure good quality, and extra time to make the product. They also include indirect costs of employee training, pilot runs, testing prototypes, designing and maintaining control systems, improvement projects, etc.
2. *Appraisal costs* are the costs of making sure the designed quality is actually achieved. They include sampling, inspecting, testing and all the other elements of quality control. They also cover administrations and audits for quality programmes.
3. *Internal failure costs*. As a product goes through the various operations in its production, it may be inspected several times. Any units that do not meet the specified quality are scrapped, returned to an earlier point in the process, or repaired. These give the internal failure cost, which is the total cost of making defective products that are detected somewhere within the process.
4. *External failure costs*. Producers normally give a guarantee with their products, and are responsible for correcting any faults. If a unit goes through the entire production process, is delivered to a customer, and is then found to be faulty, the producer must bring it back from the customer and replace, rework, or repair it as necessary. The cost of this work is part of the external failure cost, which is the total cost of making defective units that are not detected within the process, but are recognized as faulty by customers.

All things being equal, prevention and appraisal costs rise with the quality of the product, while internal and external failure costs fall. The great contribution of TQM was to show that the failure cost are high, and that organizations should spend more on prevention and appraisal to reduce the cost of failure. As the failure

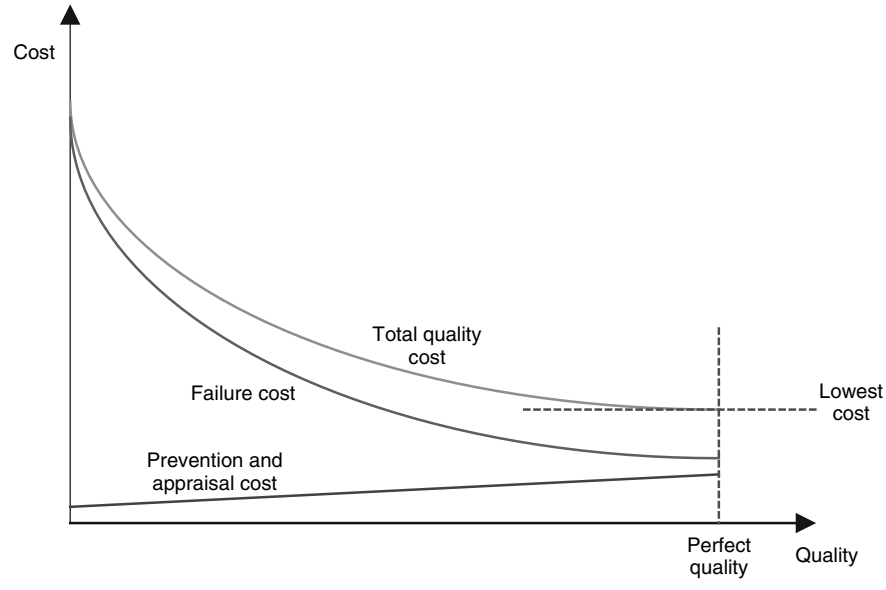


Figure 10.6 Lowest cost comes with perfect quality

costs fall with increasing quality (shown in Figure 10.6), the best way of avoiding them is to make products without defects. In other words, to make products of 'perfect quality', where every unit is guaranteed to be fault-free.

Worked example

Johnny O'Hare recorded his company's costs (in thousands of Euros a year) during a period when they introduced a major new quality management programme. From the following figures, how effective do you think the new programme has been?

Year		−3	−2	−1	0	1	2	3
Sales		2,643	2,846	2,785	3,103	3,543	3,559	3,728
Costs:	Prevention	10	12	13	45	42	42	43
	Appraisal	23	25	30	93	90	90	89
	Internal failure	113	106	110	35	33	31	29
	External failure	137	125	136	54	27	24	13

Solution

A quick look at the costs shows that they have fallen from €283,000 in year −3 to €174,000 in year 3 despite an increase of 41 per cent in sales. We can

	A	B	C	D	E	F	G	H	I
1	**Total quality management**								
2									
3	**Year**		-3	-2	-1	0	1	2	3
4	**Sales**		2643	2846	2785	3103	3543	3559	3728
5	**Costs**								
6		**Prevention**	10	12	13	45	42	42	43
7		**Appraisal**	23	25	30	93	90	90	89
8		**Internal failure**	113	106	110	35	33	31	29
9		**External failure**	137	125	136	54	27	24	13
10									
11	**Costs (% of sales)**								
12		**Prevention**	3.78	4.22	4.67	14.50	11.85	11.80	11.53
13		**Appraisal**	8.70	9.46	11.35	35.19	34.05	34.05	33.67
14		**Internal failure**	42.75	40.11	41.62	13.24	12.49	11.73	10.97
15		**External failure**	51.84	47.29	51.46	20.43	10.22	9.08	4.92
16	**Total (% of sales)**		**107.08**	**101.08**	**109.09**	**83.36**	**68.61**	**66.66**	**61.10**
17									
18									
19									
20									
21									
22									
23									
24									
25									
26									
27									
28									
29									
30									
31									
32									
33									
34									
35									

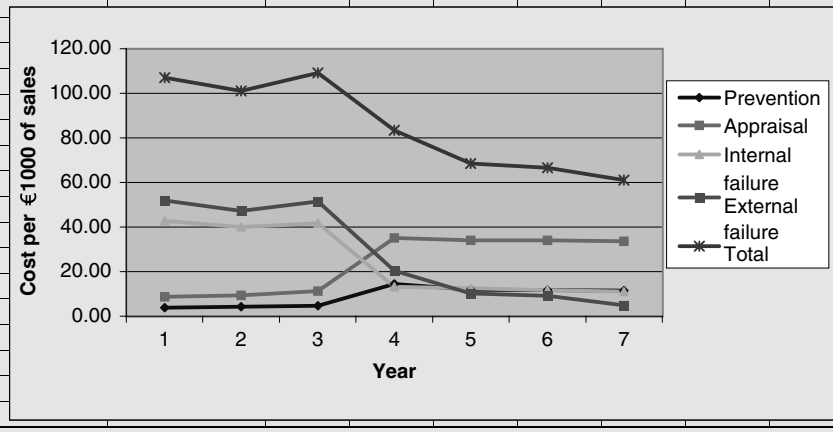


Figure 10.7 Costs of quality for worked example

make a more objective judgement by calculating the quality costs per thousand Euros of sales, as shown in Figure 10.7. The quality management programme was introduced in year zero, and this put more emphasis on prevention and appraisal, where costs have risen. As a result, product quality seems to have risen, giving lower failure costs. Overall, quality costs have fallen and sales have risen, so we must judge the programme a success.

The strength of TQM – like JIT – is that it not only sets its aims clearly, but it also shows how to achieve them. The way to achieve TQM stems from the simple observation that the longer a unit stays in a process, the more money is spent on it. If a unit becomes faulty and it then continues through its process, all subsequent work is wasted. So the cheapest way of ensuring high quality is to detect any

faults as soon as they occur. JIT takes this one step further and says that even more money can be saved by finding out why faults occur, and then taking steps to prevent them ever happening. So every operation is designed so that it cannot introduce a fault, and every person employed only passes on units that they know are of perfect quality to following operations. This approach is called *quality at source* or *jidoka*.

Traditionally, organizations used a separate quality control department to inspect the work of production departments. Quality at source stops this separation of jobs and makes everyone responsible for their own quality. This is sometimes described as 'job enlargement' as each person is now responsible for both their previous job and an inherent quality assurance job. As part of this change, people are rewarded for achieving high quality rather than traditional practices of rewarding high volumes, often with little regard to quality. This approach encourages different attitudes towards quality, including the following.

Criteria	Attitude with TQM	Traditional attitude
Importance	quality is a strategic issue	quality is a technical issue
Cost	high quality saves money	high quality costs money
Responsibility	everyone in the organization	quality assurance department
Attitude	build quality in	inspect quality in
Emphasis	preventing defects	detecting defects
Target	continuous improvement	meet specifications
Defect level	zero defects	acceptable levels of defects
Defined by	customers	the organization

One particular feature of quality at source – which is also inherent in JIT – is the search for continuous improvement. Neither is installed in one go and then left alone, but they continually evolve over time, with a series of improvements that can be absorbed without major disruption. Over time these small improvements build a momentum that can give dramatic results. Many people have give advice on ways of achieving this, notably the 'quality gurus' (Crosby, 1979; Deming, 1986; Fiegenbaum, 1983; Juran, 1988; Taguchi, 1986; Ishikawa, 1985). Some of their advice includes:

- Top managers must be committed to achieving high quality.
- The organization must adopt quality as a strategic issue.
- They must persist despite short-term problems.

- Managers must clearly state their quality objectives and what must be done to achieve these.
- Do not accept the view that there must be some defects, and refuse to accept customary levels of errors.
- Do not accept the cheapest materials, but insist on high quality.
- Do not emphasize output at the expense of quality.
- Use statistical methods to identify sources of poor quality.
- Encourage discussion of ways to improve quality (including suggestion boxes and quality circles).
- Be open to suggestions for improvement.
- Make sure everyone is properly trained to do their job.
- Do not use arbitrary numerical targets.
- Encourage a ''do it right first time'' approach.
- Search for continuous improvements.

Andon

Another feature that is common to both TQM and JIT is the recognition that any fault or problem is a sign that something has gone wrong. Rather than continuing – and probably making the problem worse – everyone has authority to stop the process and investigate. The reason for the fault is found, a solution is implemented, and procedures are changed to make sure the same fault does not occur again. This is in contrast to traditional operations that only stop production as a last resort.

Andon gives a mechanism for showing when problems are developing, and when it is time to take action. The basic format has three signals – often coloured lights – above each operation. A green signal shows that everything is working normally; an amber signal shows that operations are falling a bit behind, and a red signal shows a serious problems. The signals allow everyone to see where problems are growing, and to look for ways of solving them before they get too serious.

Respect for employees

Many companies claim that, 'our employees are our most valuable asset', but in reality they can treat employees very badly. When a company hits any financial difficulty – even a little hiccough – it inevitably reacts by laying-off employees. You might ask why they are so willing to give up their most valuable assets, and the answer is that it is often the easiest adjustment to make in the short term. JIT gives more respect to employees and really considers them to be the most valuable asset. They include in this everyone who works for the company,

removing any differences between 'managers' and 'workers'. Managers are often seen as employers who are judged by the performance of the organization, and are rewarded for high profits. On the other hand, workers are seen as employees whose wages are a drain on profits. JIT says that all employees should be concerned with the success of their organization, and they should all be treated equally and fairly. For example, all employees, regardless of their position, contribute to their organization's performance and they should be rewarded with a share of the profits.

If all employees are rewarded for the organization's performance, they are more likely actively to look for improvements. The best people to suggest improvements are those who actually work on the process, so JIT has suggestion boxes, with rewards for people offering good ideas. A more formal approach uses *quality circles*, which are informal groups of people who meet periodically to discuss ways of improving their operations. This approach is in sharp contrast to the situation in many organizations, where managers look for improvements while they work in isolation away from the process or they employ consultants who have little knowledge of the organization, its products or the process.

JIT's use of automation can also be seen as a sign of their respect for employees – although some people disagree about this. One view says that JIT encourages automation because it is more reliable and cheaper for high volume processes. Another view is that some jobs are so boring, repetitive and unsatisfying that humans should not do them if there is any alternative. Robots and computer-controlled machines can do most of the tedious work in assembly lines, and this should be automated as a matter of principle.

In return for this respect, JIT puts more demands on employees. When, for example, operators have authority to stop a process, it is assumed they will actively try to solve any problem that leads to a stoppage. They must, therefore, have the skills to identify a problem, look for the cause, find a solution and implement it. JIT also needs people who are flexible enough to do a variety of jobs. They must adopt new practices, possess relevant skills and knowledge, update their skills through training, participate actively in the running of the organization, be interested in its success, and so on.

In essence, responsibility is devolved from managers working at a distance, to people working on the shop floor. This can bring a problem with the increased stress that is put on the workforce. There is some evidence that employees who work on JIT assembly lines are subject to higher levels of stress than those who work on traditional lines. More work is needed in this area, but even a suggestion of such problems with the workforce runs contrary to JIT principles.

Summary

JIT encourages close contacts between customers and suppliers. This supports the trends towards closer integration of supply chains with organizations co-operating to obtain mutual benefits. Another important feature of JIT is its reliance on perfect quality. There has been a growing recognition that improving quality through, say, TQM can reduce costs and improve performance. JIT also requires employees

to be treated with respect. In return, it needs a skilled and flexible workforce that is dedicated to the success of their organization.

Review questions

10.9 What is JIT's view of the relationship between customers and suppliers?
10.10 Why might JIT be supplied from focused factories?
10.11 Why does an organization save money by making higher quality products?
10.12 What is the basis of quality at source?
10.13 JIT often uses an automated process, so it puts less emphasis on people. Do you think this is true?

Benefits and disadvantages of JIT

We have emphasized the role of JIT in reducing stocks of raw materials and work in progress. Some organizations have reduced these by 90 per cent (Hay, 1988). This gives a number of related benefits, such as reduced space (up to 40 per cent less), lower procurement costs (up to 15 per cent), less investment in stocks, and so on. Other benefits of JIT come from the reorganization needed to get a working system. We have already mentioned several of these, including:

- lower stocks of raw materials and work in progress;
- shorter lead times;
- shorter time needed to make a product;
- higher productivity;
- higher equipment capacity and utilization;
- simplified planning and scheduling;
- less paperwork;
- improved quality of materials and products;
- less scrap and wastage;
- better morale and participation of the workforce;
- better relations with suppliers;
- emphasis on solving problems in the process.

Unfortunately, some of these benefits can only be bought at a high price. Making high quality products with few interruptions by breakdowns, for example, can mean buying better quality, more expensive equipment. Reduced set-up times usually need more sophisticated equipment. Small batches can increases production costs. Higher skills in the workforce increase training costs and the subsequent

wage bill. Equipment must respond quickly to changing demands, so there must be more capacity. Many organizations, particularly small ones, cannot afford these costs, even when there are potential longer-term benefits.

Another disadvantage of JIT is that it is notoriously difficult to introduce, and might take many years of adjustment before the process is working properly. Then the result is somewhat inflexible, as the whole process is set up for making a specific product, and it can be difficult to change the product design, mix or production levels. This means that JIT does not work well with irregular demand, small production numbers, or specially ordered material. Seasonality also causes problems, which can be overcome in four ways. First, stocks of finished goods can be used to buffer demand; second, production can be changed to match demand; third, demand can be smoothed by pricing policies; fourth, the delivery time promised to customers can be adjusted. None of these options really satisfies JIT ideas, so real systems need some flexibility to deal with such difficulties that always exist in practice.

Some of the benefits of JIT may also be seen as disadvantages. Having frequent set-ups and small batches, for example, is essential for JIT. Most organizations find that these become a positive benefit, but some find that their costs are much higher. Similarly, JIT devolves decisions and responsibilities down to people working on the process. This kind of devolved decision making can be an advantage or a disadvantage depending on your viewpoint.

Some specific problems listed by JIT users include:

- high risks of introducing completely new systems and approaches;
- initial investment and cost of implementation;
- long time needed to get significant improvements;
- reliance on perfect quality of materials from suppliers;
- inability of suppliers to adapt to JIT methods;
- need for stable production when demand is highly variable or seasonal;
- reduced flexibility to meet specific, or changing, customer demands;
- difficulty of reducing set-up times and associated costs;
- lack of commitment within the organization;
- lack of co-operation and trust between employees;
- problems linking JIT to other information systems, such as accounts;
- need to change layout of facilities;
- increased stress in workforce;
- inability of some people to accept devolved responsibilities.

Perhaps one disadvantage of JIT is its deceptive simplicity. This has led many organizations to try JIT without understanding its underlying principles. Some

companies try to introduce elements of JIT into an existing operation. Sometimes they assume that JIT involves nothing more than reducing lead times. In an extreme case one manager circulated a note simply stating that, 'The company is introducing JIT by eliminating stocks of work in progress over the next two months. Please change your practices accordingly'. In reality JIT is an approach that needs a complete change of attitudes and operations within an organization. It is likely to take several years of careful planning and controlled implementation to get it working properly.

Summary

Just-in-time can bring considerable benefits. Many of these come from lower stock levels, but others are indirect results of the reorganization, such as higher productivity and shorter lead times. Although the benefits are clear, there can be some major disadvantages, such as inflexibility and the effort needed to get results.

Review questions

10.14 JIT eliminates stocks of work in progress. Do you think this is true?

10.15 Would it be a good idea to introduce JIT in part of a process to see how it works?

10.16 What are the three main advantage of JIT?

Extending JIT along the supply chain

Efficient consumer response

JIT forces suppliers to change the way they work, with fast deliveries, perfect quality, small batches and complete reliability. The easiest way for them to meet these requirements – which also reinforces the idea of an integrated supply chain – is to adopt JIT methods themselves. This ensures that the whole supply chain is working together with the same aims and principles. This extension of JIT along the supply chain is known by a variety of names, including *quick response*, and more commonly *efficient consumer response* (ECR).

- *Efficient consumer response* uses JIT principles to pull materials through the series of organizations in a supply chain.

Early work in ECR was done in the fashion industry. This had severe problems with its stocks, largely caused by the traditional planning of production around four seasons. At the start of, say, the summer season shops had to be full of new products in the latest styles. Shops needed high stocks to give customers a wide choice, and then wholesalers needed high stocks to quickly re-supply the shops with items that were selling well. To make sure these stocks were in place,

manufacturing was done some time before the start of the season. If demand for a product was particularly high, there would be shortages as manufacturers had already moved on to making their autumn and winter collections. If demand was low, excess supply was already being stored in shops and wholesalers. At the end of each season there were major sales to get rid of the remaining less popular items, and major restocking in preparation for the next season.

The industry realized that it could get enormous savings if it smoothed out enormous operations. The way to do this is not to have huge stocks sitting in the supply chain, but to move units quickly, and respond to customer demands by more flexible manufacturing. Now they use just-in-time operations and link information systems so that they can 'pull' materials through the supply chain. When a retailer sells a unit, their cash register automatically sends a message to the wholesaler requesting a replacement. In turn, the wholesaler's system sends a message to the manufacturer asking for a delivery. The manufacturer is not bogged down in making excessive amounts of items that are later sold at discounts, but responds quickly to the demand and replaces garments that have actually been sold.

ECR extends the benefits of JIT to the whole supply chain. So it brings lower stocks, better customer service, lower costs, more responsive operations, improved space utilization, less paperwork, and so on. Organizations introducing ECR in the 1990s reported a string of benefits. Quaker Oats, for example, reported a threefold increase in stock turnover, 65 per cent lower stocks and 77 per cent reduction in paperwork (Boden, 1995). Integrated Systems Solutions reported 3–4 per cent increase of service level, 40–50 per cent reduction in stock and 2–3 times increase of stock turnover (Margulis, 1995).

Features of ECR

It is not necessarily physical transport that slows the flow of materials through a supply chain, but the associated flow of information. It might take a month for an organization to prepare the details for a purchase, collect information, send orders, arrange payments, etc. while delivery only takes a day. So ECR only became feasible when a practical method of control was designed. With JIT this came with *kanbans*; with ECR it came with EDI (electronic data interchange). Each organization's control system sends a message to suppliers and signals the need for more materials using an 'electronic *kanban*'. Some systems go further and hand over more responsibility to the supplier in *vendor-managed inventory*. Then the supplier becomes responsible for maintaining stocks at their customers' operations, checking the availability, organizing deliveries and all other aspects of inventory control that make sure stocks are available when needed.

Like JIT in general, ECR is a deceptively simple idea, but it needs major changes to operations and can only be used in certain circumstances. If the supply chain starts with potatoes, they are grown in a particular season and farmers cannot suddenly grow a crop at short notice. Another problem comes with the length of the supply chain, as a single organization that does not want to be involved – or cannot adapt – will disrupt the flow. If the supply chain crosses a slow international border, or includes an area where productivity is low, or hits other problems, the delays become unacceptable and ECR cannot work.

Introducing JIT can be a huge undertaking, fundamentally changing the way an organization works. When JIT is extended to ECR, implementation becomes an even bigger issue. This is probably why organizations have seemed slow to introduce it. By 1997 almost no organizations had a fully established, working ERC system (P-E Consulting, 1997).

Summary

The principles of JIT can be extended to organizations along the supply chain. Methods like ECR give an integrated flow of materials that is pulled by final customer demand. This can bring considerable benefits, but it needs a lot of effort for implementation.

Review questions

10.17 The benefits of ECR are so clear that all organizations should be moving in this direction. Do you think this is true?

10.18 What is an 'electronic *kanban*'?

Comparisons with other methods of inventory management

Differences in approach

At first sight JIT seems radically different from the other methods of inventory control. For example, independent demand methods – and to a lesser extent MRP – work within an existing organization, while JIT says that major changes are needed before it can start working. Other methods assume that fixed values for costs, demand, lead time, etc., while JIT assumes these can be altered and positively looks for improvements. Other methods look for the best way of controlling stocks, while JIT actively works to eliminate them. Other methods, particularly MRP, need a lot of computer support, while JIT simplifies operations so that computers are not needed for day-to-day operations.

Despite the significant differences in viewpoint, JIT has several points in common with other approaches to inventory management. They all accept that costs can be high and look for savings; they know that stocks affect all other aspects of operations; they recognize that stocks should be low, provided that service levels can be maintained, and so on. But there are more basic similarities. Independent demand methods often track stock levels and place an order of fixed size whenever this declines to the reorder level. Some people say that this is exactly what *kanbans* do, placing an order for a container full whenever stock declines to a predetermined level. Then *kanbans* do not give fundamentally new methods, but only provide a mechanism for running traditional methods. With this view, the changes brought about by JIT – short lead times, frequent orders, small quantities, perfect quality, etc. – are results that good managers should already be aiming for.

The opposite view says that there are fundamental differences between JIT and other approaches. We can, for example, list the following differences between JIT and MRP.

- JIT is a manual system, while MRP relies on computers.
- JIT 'pulls' materials through a process, while MRP 'pushes' them with predefined schedules.
- JIT emphasizes the control of operations, while MRP is more concerned with planning.
- JIT emphasizes physical operations, while MRP is largely an information system.
- JIT puts control of the process on the shop floor, while MRP gives control to more distant planners.
- JIT works with a minimum amount of data, while MRP tries to collect all possible data.
- JIT reduces the amount of clerical effort, while MRP increases it.
- JIT needs a constant rate of production, while MRP can work with varying production.
- JIT has reducing set-up cost as a priority, while MRP considers this to be fixed.
- JIT can be easily understood by everyone using it, while MRP is more difficult to understand.
- JIT reduces batch sizes to a minimum, while MRP uses batching rules to increase them.
- JIT typically carries hours' stock of material, while MRP carries stocks for days.

Working together

Despite their differences, every approach to inventory management has similar overriding aims of giving a smooth flow of materials, high customer service and low costs. Sometimes these aims are best achieved by not using one particular method, but by using the best combination of methods for specific purposes. It might be, for example, that an organization has some operations that work better with MRP (perhaps lower volume, batch processes), others that work better with JIT (perhaps higher volume mass processes) and others that work better with independent demand methods (perhaps retail operations). The different approaches can often give similar thinking, as we mentioned above with the similarity between *kanban* and independent demand methods. We have also mentioned that independent demand models can work within MRP, for setting batch sizes and delivery times. Perhaps, then, we should not consider these methods in isolation, but should be looking for the best overall approach for particular circumstances.

Some people say that MRP is more a planning system, while JIT is more a control system. One useful approach, then, is to use a combination of both systems. MRP does the overall planning, making sure that in the longer term there are enough materials arriving to support the process, while JIT controls the flow of materials within the process. You might imagine MRP controlling the external flow of materials into the organization, with JIT taking over the internal flow within the organization (as shown in Figure 10.8).

Obviously the choice of approach must depend on particular circumstances and be a decision for management. Generally, organizations making lower volumes of more varied products in batches might do best using MRP; those making higher volumes of similar products in continuous operations might be better using JIT; and those in between might use some combination of the two (Voss and Harrison, 1987). Both MRP and JIT only work well in specific circumstances,

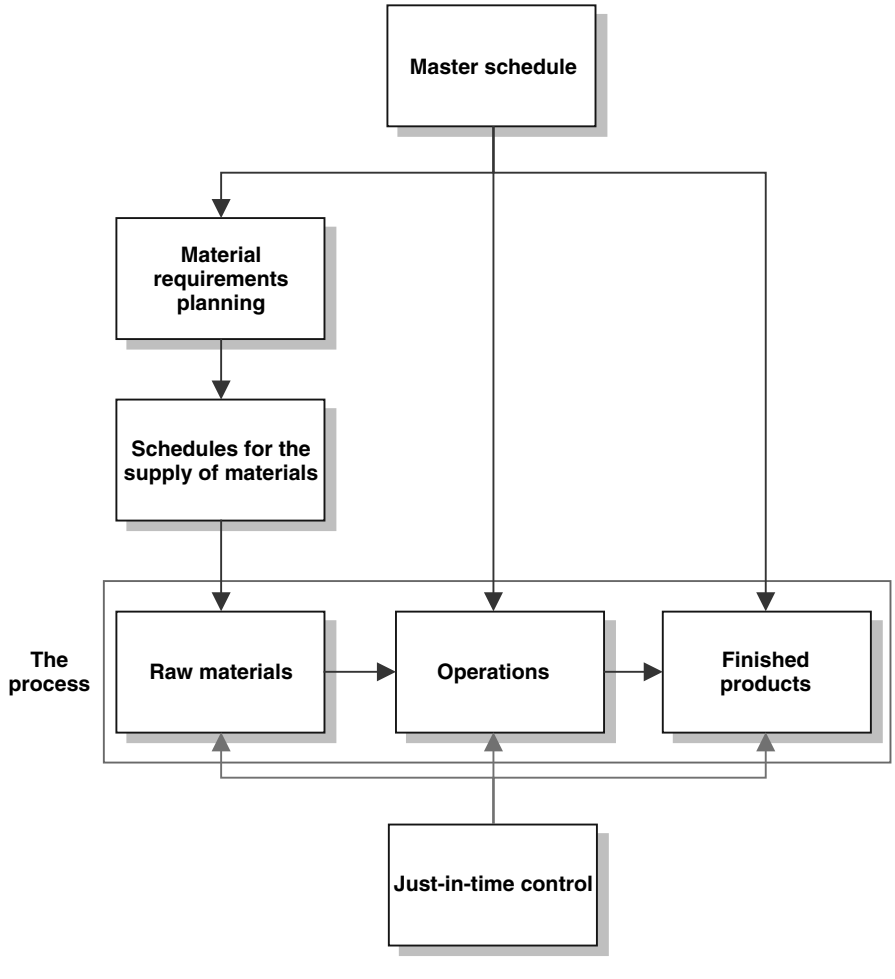


Figure 10.8 Combining MRP and JIT

and implementation of any such system can be a major undertaking; in contrast, independent demand models are flexible, and give good results for a wide variety of circumstances. If circumstances are very complicated an organization might use project management methods, such as PERT, and this opens entirely new opportunities. The following table shows the area where each method is most commonly used.

Approach	Typical use
Independent demand methods	
deterministic models	a wide variety, but particularly finished goods and work in progress where there little uncertainty
probabilistic models	a wide variety, but particularly finished goods and work in progress where there is some uncertainty
Dependent demand methods	
material requirements planning	raw materials and work in progress for batch production
just-in-time	raw materials and work in progress for high volume, continuous production

Summary

Despite their obvious differences, independent demand methods, MRP and JIT share some points in common. The choice of best depends on specific circumstances, particularly the type of process. Each need not work in isolation, but they can work together to give the best overall result.

Review questions

10.19 What are the two most important factors in choosing a method of inventory management?

10.20 The alternative approaches to inventory management are so different that each has to be used in isolation. Do you think this is true?

Chapter review

- This chapter described just-in-time methods, which organize operations so that they occur at exactly the time they are needed.
- JIT is based on an aim of eliminating all waste. Stocks are seen as a waste of resources that can be eliminated by properly co-ordinating the supply and demand. The characteristic approach is to identify problems and solve them, rather than hiding them under excessive stocks.

- JIT needs many changes to the way that organizations view their operations. Even when they are willing to make these changes, JIT does not work well in all circumstances. It is generally most successful with large-scale assembly.
- JIT needs a simple means of controlling the supply of materials. This is usually done by *kanbans* which signal a demand pull from later operations to earlier ones.
- JIT encourages close contacts between customers and suppliers. This supports the trends towards closer integration of supply chains with organizations co-operating to obtain mutual benefits.
- Another important feature of JIT is its reliance on perfect quality. There has been a growing recognition that improving quality through, say, TQM can reduce costs and improve performance.
- JIT also requires employees to be treated with respect. In return, it needs a skilled and flexible workforce that is dedicated to the success of their organization.
- Just-in-time can bring considerable benefits. Many of these come from lower stock levels, but others are indirect results of the reorganization. There are also some disadvantages with JIT, such as inflexibility and the effort needed to get results.
- The principles of JIT can be extended to organizations along the supply chain. Methods like ECR give an integrated flow of materials that is pulled by final customer demand.
- Despite their obvious differences, independent demand methods, MRP and JIT share some points in common. The choice of best depends on specific circumstances, particularly the type of process. Each need not work in isolation, but they can work together to give the best overall result.

Project

The aim of this project is to compare the costs of independent demand methods, material requirements planning and just-in-time. For this you have to design a simulation – or some other method – for finding the cost of using each method in an organization. You might start with a specific problem – perhaps a small batch production – and see which method gives lowest costs. You can find the factors that have most effect on costs and how these can be controlled. Then by changing the features of the model, you can find the method that works best in different circumstances.

Problems

10.1 The output from an assembly line is 1,600 units in an 8-hour shift. The item is moved in containers, each of which holds 50 units. Each container spends an

average of 30 minutes in the production part of a cycle, and 20 minutes in the demand part. How many *kanbans* should be used for the item? What is the stock of work in progress? A scheme has been proposed to reduce the time in the demand part of the cycle to 15 minutes. How many *kanbans* should be used now? What is the margin of safety in each case?

10.2 Three years ago JR Martingale Ltd introduced a new quality assurance programme. The costs (in thousands of pounds) over the past seven years have been as follows;

Year	−3	−2	−1	0	1	2	3
Sales	1,976	2,019	1,905	2,374	2,770	2,841	2,689
Prevention	14.3	15.4	15.7	52.3	57.9	61.2	60.3
Appraisal	16.5	18.6	19.3	89.5	73.6	71.4	75.5
Internal failure	108.3	113.5	121.8	42.6	38.5	41.7	36.4
External failure	106.3	127.4	45.7	34.6	33.1	30.9	26.2

Has the quality assurance programme been a success?

Discussion questions

10.1 If you were in hospital and needing a blood transfusion, would you rather the transfusion service used an independent demand method of controlling stocks of blood, or a just-in-time system? What does your answer tell you about JIT in other organizations?

10.2 JIT is not a radical new approach, but formalizes some of the effects that good managers should be aiming for (such as frequent deliveries, short lead times, low reorder costs, etc.). What do you think of this view?

10.3 Because they cannot hold stocks of products, services must use some form of JIT. How do they do this, and why did it take so long for manufacturers to recognize the benefits?

10.4 JIT reduces waste in an organization. What is meant by waste in this sense, and how can it occur? If a manufacturer only has enough demand to work for 7 hours in an 8-hour shift, would it be more wasteful to leave all operations idle for an hour, or to make extra units and put them into stock?

10.5 If higher quality reduces costs, why does a high quality car (or apparently anything else) cost more than a low quality one?

10.6 JIT aims for simple operations; MRP designs huge, complicated systems. What are the main differences and similarities between these two approaches? Are they so different that they can never work together?

10.7 Managers do different jobs to workers on the shop floor, so they should be rewarded differently. Managers look after the interest of shareholders, so they should be given shares and bonuses based on profits: workers make the products, so they should be given a bonus for high productivity. What do you think of such arguments?

References and further reading

Bicheno, J. (1999) Kanban: the old and the new, *Control*, September, pp. 22–26.

Boden, J. (1995) A movable feast, *Materials Management and Distribution*, November, pp. 23–26.

Crosby, P.B. (1979) *Quality is Free*, New York: McGraw-Hill.

Deming, W.E. (1986) *Out of the Crisis*, Cambridge, MA: MIT Press.

Fiegenbaum, A. (1983) *Total Quality Control*, New York: McGraw-Hill.

Fucini, J. and Fucini, S. (1990) *Working for the Japanese*, New York: Free Press.

Hay, E.J. (1988) *The Just-in-Time Breakthrough*, New York: John Wiley.

Hunter, A., King, R. and Lowson, B. (1999) *Quick Response*, Chichester: John Wiley.

Ishikawa, K. (1985) *What is Total Quality Control?*, Englewood Cliffs, NJ: Prentice Hall.

Jacobson, G. and Hillkirk, J. (1986) *Xerox: American Samurai*, New York: Macmillan.

Juran, J.M. (1988) *Juran on Planning for Quality*, New York: Free Press.

Louis, R. (1997) *Integrating Kanban with MRP II*, Cambridge, MA: Productivity Press.

Margulis, R.A. (1995) Grocers enter the era of ECR, *Materials Management and Distribution*, February, pp. 32–33.

Monden, Y. (1994) *Toyota Production System* (2nd edition), London: Chapman and Hall.

Ohno, T. and Mito, S. (1988) *Just-in-Time for Today and Tomorrow*, Cambridge, MA: Productivity Press.

P-E Consulting (1997) *Efficient Customer Response*, Surrey: P-E Consulting/Institute of Logistics.

Sewel, G. and Wilkinson, D. (1995) *Integrating JIT, MIS and TQM*, London: Butterworth-Heinemann.

Shingo, S. (1981) *Study of Toyota Production System from Industrial Engineering Viewpoint*, Tokyo: Japanese Management Association.

Taguchi, G. (1986) *Introduction to Quality Engineering*, Tokyo: Asian Productivity Association.

Voss, C.A. and Harrison, A. (1987) Strategies for implementing JIT, in C.A., Voss (ed.), *Just-in-Time Manufacture*, Berlin: IFS/Springer-Verlag.

Womack, J., Jones, D.T and Roos, D. (1990) *The Machine that Changed the World*, New York: Rawson.

Solutions to Review Questions

1 Stocks and inventories

1.1 An inventory is the list of items stored, while stocks are the actual units. (This distinction is becoming less clear.)

1.2 A single copy of Everyperson's Encyclopaedia.

1.3 In this context a supplier is anyone or anything that adds units to stock, while a customer is anyone or anything that removes units.

1.4 They decouple adjacent operations, so that each can continue if there are problems with the other.

1.5 No – there are other reasons for holding stock.

1.6 Depending on the organization, lubricating oil might be a raw material, work in progress, finished goods, consumable or spare parts. It is most commonly a consumable.

1.7 The series of activities and organizations that a product moves through from initial suppliers through to final customers.

1.8 To some extent – but there are other ways of getting good customer service.

1.9 Because they have a common aim of supplying products that give final customer satisfaction.

1.10 No – it can improve operations, reduce its costs, increase sales, reduce risk, and so on.

1.11 There have always been adjustments to inventory management, but perhaps the key factors encouraging rapid change now are fiercer competition and improving technology.

1.12 No, they are all inter-connected.

1.13 It seems likely for the foreseeable future.

1.14 Because it is expensive to hold stock, and organizations are trying to reduce costs.

1.15 No.

1.16 Because production falls before demand, and extra demand is met by reducing stocks and liquidating companies.

1.17 Probably, but there seems no practical way of achieving this.

2 Stock within an organization

2.1 Procurement, inward transport, receiving, warehousing, stock control, order picking, material handling, outward transport, recycling, location, communication and a range of other activities.

2.2 No.

2.3 Some decisions are made at all levels – strategic, tactical and operations.

2.4 No.

2.5 A lean strategy puts more emphasis of removing waste, while an agile strategy puts more emphasis on customer service.

2.6 Yes.

2.7 Broader strategic decisions, the external environment and organizational strengths.

2.8 The strategies may demand higher levels of performance from the stock holdings than the organization can provide.

2.9 Generally, yes.

2.10 Almost completely.

2.11 A rough estimate has the value of food in a freezer as £250 and storage costing 20 per cent of this a year. Costs are probably around £1 a week.

2.12 A high one.

2.13 Usually monitoring improvements over time, but sometimes comparing performance with other organizations.

2.14 Because it is fixed by the timing and size of orders.

2.15 Periodic reviews, fixed order quantity and matching demand.

2.16 Usually from forecasts based on historical demand.

3 Economic order quantity

3.1 The main assumptions are: a single item is considered, demand is known and is constant and continuous, lead time is zero and replenishment is instantaneous, unit, reorder and holding costs are all known exactly and are fixed, no shortages are allowed. There are a number of other assumptions implicit in the analysis.

3.2 The economic order quantity is the order size that minimizes total inventory costs for a simple inventory system.

3.3 Because average stock levels are higher, giving high holding costs. The increase in holding cost component is greater than the decrease in reorder cost component.

3.4 The order size.

3.5 Neither – at the EOQ the reorder cost component equals the holding cost component.

3.6 The EOQ identifies the minimum cost per unit time, so orders of any other size will increase the costs.

3.7 No – the cost rises slowly.

3.8 No.

3.9 No – the calculated value for EOQ is unlikely to be exactly optimal, but the analysis gives useful guidelines and good results in many different circumstances.

3.10 Without any further information, we can guess that the variable cost curve will rise less steeply for orders larger than EOQ, so it is generally better to use larger orders.

3.11 The amount of stock on hand when it is time to place another order.

3.12 The stock cycle is 3 weeks long, so the lead time is between two and three stock cycles. There will still be two orders outstanding when it is time to place another order (so ROL $= 7 \times 10 - 2 \times 30 = 10$ units).

3.13 Yes – unless there is an alternative way of signalling that stock level has declined to the reorder level, such as a two-bin system.

4 Models for known demand

4.1 Depending on circumstances, any costs can vary with order quantity.

4.2 A calculated economic order quantity that can actually be used (i.e. is on the valid cost curve).

4.3 Either at a valid minimum or at a cost break.

4.4 Not necessarily – the best order size might be at a break point.

4.5 Zero.

4.6 It can be either greater than or less than, depending on the relative values of variables.

4.7 All things being equal, finite production rates lead to lower stock levels, lower holding costs and larger batches.

4.8 The calculated values become imaginary – which is why the analysis is only valid when the replenishment rate is greater than demand.

4.9 When there are no stocks of the item, and customers are willing to wait until the next delivery.

4.10 The obvious problem is defining a reasonable shortage cost. This is almost inevitably a matter of subjective evaluation and agreement.

4.11 No – it depends on the relative costs.

4.12 When customers demand an item that is out of stock and they do not wait for a back-order.

4.13 The costs exactly match gross revenue for this item, so it does not matter whether it is stocked or not.

4.14 Holding stock always involves costs, so if shortages are allowed the costs may be minimized by not holding any stock at all.

4.15 If the assumptions remain the same it does not change.

4.16 No – it is the demands that are independent not the items.

4.17 By adding an additional cost for storage space. This increases the effective holding cost, and reduces the order quantity and average stock level.

4.18 Less – as the constraint lowers the average stock level.

4.19 The demand comes in integer amounts every period, and varies from one period to the next.

4.20 Not quite – D_i is the demand in the ith period of a cycle.

4.21 Because it is based on a number of assumptions and approximations.

5 Models for uncertain demand

5.1 In any area of supply and demand, but perhaps most importantly in demand, costs, lead time and supplied reliability.

5.2 No – deterministic models give useful results in a wide range of circumstances, particularly when the amount of uncertainty is small.

5.3 Variability means values change over time: the changes may be known with certainty or be uncertain. Uncertainty means that values follow a known probability distribution. Ignorance means we do not know anything about the values.

5.4 Because we can allow for any variation outside the lead time by the timing and size of the next order. When we are inside the lead time it is too late to make any adjustment.

5.5 A model where an item is only held in stock for a single period, and any unsold units are scrapped at the end of the period.

5.6 These models are most useful for highly seasonal goods, such as Christmas fare, Easter eggs, winter clothes, and so on.

5.7 All of the first Q units have an expected profit, but the $(Q + 1)$th unit has an expected loss that reduces the overall profit.

5.8 Any values might be difficult to find, but shortage cost are the most consistently difficult.

5.9 Stocks of spare parts are usually held at a specified level, and any units used are replaced immediately so the stock level is more relevant than the order quantity.

5.10 Raise it – thereby avoiding the shortage costs.

5.11 When there is wide variation in lead time demand and when the shortage cost is relatively low.

5.12 An amount of stock that is held in reserve to cover unexpectedly high demand or late deliveries.

5.13 This measures the service given to customers. There are several related measures, but we use 'cycle service level' which is the probability there are no shortages in a stock cycle.

5.14 By holding more safety stock.

5.15 If there is a lot of variability in lead time demand, very high safety stocks – in principle infinite – would be needed to give 100 per cent service level.

5.16 Because any uncertainty outside the lead time can be allowed for by the size and timing of the next order. During the lead time it is too late for such adjustments.

5.17 Increase.

5.18 No – the safety stock should depend on the variability of demand and not its absolute value.

5.19 Size and variability of lead time demand, and desired service level.

5.20 No.

5.21 By subtracting the amount of stock on hand and on order from the target stock level.

5.22 A periodic review method (all other things being equal).

5.23 There are several possible reasons, including ease of administration, setting a routine for purchasing, ease of combining orders, and so on.

6 Sources of information

6.1 The notional part of a MIS that includes the procedures, systems and other resources that ensure everybody in an organization has the information about stocks that they need.

6.2 In principle a system could work with current stocks, reorder levels, reorder quantities and goods on order. In practice much more information is needed.

6.3 There are many of these, including data validation, recording stock movements, automatic purchasing, invoice clearing, and so on.

6.4 To all parts of the organization, to linked organizations (principally suppliers and customers), and to any other place where the information is needed.

6.5 Up to a point. The quality of decisions depends on the skills of managers, but they need enough information to make informed decisions. There also comes a point where they have enough useful information and any more gives an overload.

6.6 Yes.

6.7 There is not really an absolute value (as 'stock is only worth what someone will pay for it'), so we use some kind of estimate.

6.8 LIFO.

6.9 To see which items ones should be given most – and least – attention.

6.10 B items.

6.11 Because procurement initiates the movement of materials into an organization, and hence into the stocks.

6.12 Not really – procurement processes the information and initiates the move, but does not physically move the materials (which is probably done by a Transport Department).

6.13 Not until the very long term – possibly never.

6.14 Inventory control is a management function that makes decisions about stock, while warehousing physically stores and looks after the stock.

6.15 There are many of these including location, size, equipment used and layout.

6.16 Manual, mechanized and automated.

6.17 It identifies the product, gives basic information, protects the contents, makes handling easier and helps market and promote the product.

7 Forecasting demand

7.1 All current plans and decisions become effective at some point in the future. They need information about prevailing circumstances, and this must be forecast.

7.2 Judgemental, causal and projective forecasting.

7.3 Many factors, including the forecast purpose, availability of quantitative data, time in the future covered, relevant external factors, cost, effect of errors, time available, and so on.

7.4 No.

7.5 When there is no historical data for a quantitative forecast.

7.6 Personal insight, panel consensus, market surveys, historical analogy and the Delphi method.

7.7 Unreliability, conflicting views from experts, cost of data collection, lack of available expertise, and so on.

7.8 A series of observations at regular intervals.

7.9 Forecasts find the underlying pattern in demand, but they cannot deal with short-term, random noise. Errors are also introduced by not finding the true underlying pattern and unexpected changes in the demand.

7.10 Positive and negative errors cancel each other, so the mean error should have a value around zero unless there is bias.

7.11 By calculating the errors with each method. All things being equal, the best method is the one with smallest errors.

7.12 It looks for relationships between variables, and then forecasts the value of one (the dependent variable) from the known value of the other (the independent variable).

7.13 A method of finding the line of best fit (i.e. the line that minimizes the sum of squared errors) through a set of points.

7.14 The proportion of the sum of squared errors from the mean that is explained by the regression.

7.15 No – it only measures the strength of the relationship, not a cause and effect

7.16 Because older data tends to swamp more recent (and more relevant) data.

7.17 By taking an average over fewer periods.

7.18 Because the weight given to data declines exponentially with its age, and the method smoothes the effects of noise.

7.19 By using a higher value of α.

7.20 Because it cannot get the correct timing of seasons and lags behind trends.

7.21 By using the method described but setting all the seasonal indices to 1.0.

7.22 We might be tempted to over-estimate demand, but the answer really depends of the relative consequences of too high and too low stocks. The real answer is that it is better to produce an accurate forecast.

7.23 Maybe – but the details of the method are probably less important than the planning and preparation.

8 Planning and stocks

8.1 To some extent – but under different circumstances any level of decision might be considered most important.

8.2 By the strategic plans which provide the context for lower level decisions and by tactical plans in related areas.

8.3 The overall capacity sets the maximum throughput – the aim of inventory managers is to have enough capacity to meet forecast demand, so that stock is not a bottleneck in broader operations.

8.4 Because it is unlikely that the initial plan will satisfy all constraints, objectives and everybody concerned.

8.5 It wouldn't – aggregate plans do not deal with individual products.

8.6 A schedule of production for each family of items, typically by month and covering the next few months.

8.7 Because they allow a buffer which means that production does not have to match demand exactly.

8.8 It is relatively straightforward, easy to understand, convenient, fast to compare alternatives, etc.

8.9 One way has a cumulative supply line that is close to a cumulative demand line, and with few changes in gradient.

8.10 When small deviations from optimal solutions give significant increases in cost or some other penalty.

8.11 To disaggregate the aggregate plan and show a timetable for making individual items.

8.12 These mainly come from the aggregate plan and include available capacity, actual customer orders, costs and any other specific constraints.

8.13 To produce timetables for individual pieces of equipment and other resources.

8.14 Because scheduling rules generally give good results with little effort, and do not require as much data, models, time or some other requirement of more rigorous methods.

8.15 To monitor progress, to make sure that planned schedules are actually being achieved, to warn of problems, to make minor adjustments to schedules, to give feedback, and so on.

8.16 Simulation uses a dynamic representation of a system to replicate its operation over time. It can be used to find the effects of planned changes.

8.17 Because it allows the effects of proposed changes to be examined quickly and in detail, without affecting the performance of real operations.

9 Material requirements planning

9.1 Independent demand methods assume that all demands are unrelated and forecast demand from historical figures: dependent demand methods assume there is some link between demands and typically forecast demand from production plans.

9.2 No. They are widely used, but in some circumstances do not perform as well as dependent demand methods.

9.3 A master schedule, the bill of materials, current stocks, lead times – and any other relevant information.

9.4 By calculating the gross requirement and then subtracting free stock and scheduled receipts.

9.5 No – MRP was initially developed for manufacturing, but has since been applied in many different circumstances.

9.6 The direct linking of supply to demand reduces stocks and gives a number of related benefits.

9.7 The requirements that limit its applicability, the complexity of the systems and inflexibility.

9.8 Timetables of orders, changes to orders, timetables for operations, exception reports, performance reports, planning reports, inventory transactions, etc.

9.9 Because of the difficulty of finding reliable forecasts from historical data.

9.10 No.

9.11 A rule to suggest how many separate smaller orders should be combined into a single larger order.

9.12 Feedback showing the consequences of the MRP schedules – probably followed by adjustments to schedules and capacity.

9.13 There are many ways ranging from basic adjustments of the calculations through to major extensions into different operations, functions, organizations, etc.

9.14 Manufacturing resource planning, which extends the MRP approach to a wide range of functions.

9.15 Yes – in common with all other approaches of this type.

10 Just-in-time

10.1 Operations occur just as they are needed, so materials are delivered just as they are needed, and stock is eliminated.

10.2 Other methods assume that stocks are essential and look for ways of minimizing the costs: JIT assumes that stocks are a waste of resources and looks for ways of solving the problems that are hidden by stock.

10.3 Possibly – but usually there are many other benefits.

10.4 Generally high volume continuous processes, often using automated assembly lines.

10.5 No – there are many examples of JIT working very successfully in small organizations and services; some hamburger restaurants, for example, come close to JIT.

10.6 Push systems complete planned operations and then push materials on to the next operations; pull systems wait until a message is passed backwards, and then they return the materials requested.

10.7 To control the flow of materials in a JIT system.

10.8 Materials are only moved in containers, each of which has a *kanban*, so the number of *kanbans* sets the number of containers and hence the amount of work in progress.

10.9 They are long-term partners who co-operate and have a mutually beneficial trading arrangement.

10.10 Because focused factories specialize in making one item very efficiently, using specialized equipment, long production runs and other features encouraged by JIT.

10.11 Because savings on failure costs more than cover extra costs for prevention and appraisal.

10.12 Each person in a process is responsible for passing on units they know to have perfect quality.

10.13 No.

10.14 Not really – JIT minimizes stocks of work in progress but does not eliminate them.

10.15 JIT needs a fundamental change in attitudes, plans, operations and procedures. It is not the sort of thing that can be tried as a small experiment.

10.16 There are many advantages including reduced stock, easier planning, higher quality, better control, lower costs, etc.

10.17 Not really – like JIT itself, ECR only works well in certain circumstances.

10.18 A request passed back through the supply chain by EDI to supply more materials.

10.19 There is some debate about this, but probably the type of process and the stock held.

10.20 No – the methods can work together.

Index

ABC analysis 207–11
accounting information 202–10
  ABC analysis 207–11
  return on assets 45–8
  trading accounts 202–3
  value of stocks 48–51, 204–7
  see also costs
actual cost of stock 48–51, 204
aggregate plans 274–85
  aims 275–6
  designing 278–80
  master schedule 274
  spreadsheets 276–8, 282–4
  tactical plans 274
aggregate stocks 24–8
agile operations 41
aims of
  dependent demand methods 57–8, 307–8
  forecasting 232
  independent demand methods 57–8
  inventory management 36–44
  just-in-time 341–4
  linear regression 241–2
  MRP 310–11
  procurement 211–12
  simulation 292–3
  warehousing 218
alliances 18
amplification of demand 15–17
approaches to inventory control 55–8
assumptions of the economic order quantity 67–8
automation in warehouses 222–3
average investment constraint 132–4

B2B 217
B2C 217
back-orders 120–8, 276
basic questions of inventory control 55–8
batch size *see* order quantity
batching rules for MRP 327–9
bill of materials 311–21
bottlenecks 270
breadth of supply chain 12–14
buffer 7–9
business cycle 26–8
business strategy 37–44

capacity planning 270–3
  bottlenecks 270
  methods 271–3
capacity requirements planning 330
causal forecasting 233–4, 241–8
  linear regression 241–8
  multiple regression 247–8
  non-linear regression 242
  spreadsheets 243, 247
changing views of stock 24
classification of
  costs 52–3
  forecasting methods 233–4
  inventory management methods 57–8
  stock 9–10
closed loop MRP 330–1
coefficient of correlation 245–7
coefficient of determination 244
co-managed inventory 22
communications
  improving 20–1
  in logistics 33
concentration of ownership 21
constraints 128–34
  on average investment 132–4
  on storage space 129–32
consumables 9
continuous improvement 360
control of schedules 291–2
co-operation in supply chain 15–19
  achieving 17–19
  benefits 15–19
  problems with 15–17

corporate strategy 37–44
costs
  classification 52–3
  delivery 109–12
  fixed 71
  holding 52–3, 69
  implied 87–9
  inventory 48–53
  reorder 52, 69
  shortage 53, 69, 122, 125, 161
  types of cost 52–3
  uncertain 87
  unit 52, 69, 100–12
  value of stock 48–51
  variable 71–3, 100
cost leadership 40
cross-docking 22–3, 219
customer
  definition 7
  importance 10
  service 21
  types of 9–10
cycle
  business 26–8
  counting 201
  procurement 213–15
  stock 9
  time 69, 72
cycle service level 172–3

delivery cost 109–12
Delphi method 236–7
demand
  amplification 15–17
  dependent 57, 305–73
  discrete, variable 135–41, 154–65
  forecasting 211, 229–66
  independent 57–8, 305–73
  intermittent 164–5
  known 99–144
  lead time 90–4, 149, 173–81
  Normally distributed 154–5, 174–8, 180–1, 183–5
  uncertain 86–7, 147–89
  varying 15–17
dependent demand
  definition 57–8, 307–8
  just-in-time 341–73
  material requirements planning 307–40
  methods 305–73
derivation of the economic order quantity 69–72
design
  aggregate plans 278–80
  master schedule 286–7
  operational schedules 288–92
  supply chain 13–15
  warehouse 220–3
deterministic models 99–144
  see also known demand
discounts in unit cost 100–12
  spreadsheet for calculations 108–9
discrete values for EOQ 78, 82–6
discrete variable demand
  known demand 135–41
  uncertain demand 154–70
distribution resource planning 334
drop shipping 22

economic order quantity 65–98
  adjusting 87–8
  assumptions 67–8
  background 65–8
  calculations 72–7
  definition 72
  derivation 69–72
  discrete values 78, 82–6
  rounding to integers 78, 82–6
  sensitivity 78–82
  spreadsheet for calculations 75–7
  variables used 69
  weaknesses 77–8
ECR 365–6
EDI 20
efficient customer response 365–6
EFT 20
electronic data interchange 20
electronic fund transfer 20
electronic point of sales 20
EOQ *see* economic order quantity
EPOS 20
e-procurement 20, 216–17
e-purchasing 20, 216–17
environmental concern 23–4
enterprise resource planning 334–5
ERP 334–5
errors
  in forecasts 239–41
  in parameters 86–7
exponential smoothing 252–5

FIFO 49–51, 204
finished goods 9
finite replenishment rate 113–20
  spreadsheet for calculations 117–19
first-in-first-out 49–51, 204
fixed cost 71
fixed order quantity methods 181
  advantages 185–6
  see also economic order quantity

flow diagrams 293
focus of strategies 41–2
forecasting 211, 229–66
  aims 232
  causal 233–4, 241–8
  choice of method 232
  classification of methods 233–4
  errors 239–41
  exponential smoothing 252–5
  judgemental 233–7
  methods 230–3
  moving averages 249–52
  noise 239
  planning 260–2
  projective 233–4, 248–60
  sensitivity 250, 254
  seasonality and trend 256–9
  simple averages 248–9
  spreadsheets 251, 254–5, 257–8
  time horizons 233
  time series 237–41
  tracking signal 254–5
functional strategies 37

gross stock 92–4

historical analogy 236
holding cost 52–3, 69

implementing strategies 42
implied costs 87–9
independent demand
  definition 57–8, 307–8
  information needed 193–302
  models 63–192
  problems 307–9
  see also models for
information 193–302
  ABC analysis 207–11
  accounting 202–10
  basic 202–4
  cycle counting 201
  inventory management information system 195–201
  management information system 196
  needed 196–8
  procurement 211–18
  sources 195–227
  stocktaking 200–1
  supply and demand 211–17
  transaction recording 198–203
  value of stock 48–51, 204–7
integer values for the economic order quantity 78, 82–6
integrated logistics 34–6
intermittent demand 164–5
invalid minimum 102–12
inventory
  definition 4
  see also stock
inventory control *see* inventory management
inventory management
  aims 36–44
  approaches 55–8
  as part of supply chain management 11–12
  basic questions 55–8
  co-managed 22
  definition 7
  dependent demand 57–8, 305–73
  independent demand 57–8, 63–192
  information needed 193–302
  information system 195–201
  levels of decision 36–8
  models *see* models for
item
  coding 20
  definition 4

jidoka 356–63
JIT *see* just-in-time
judgemental forecasts 233–7
  methods 235–7
just-in-time 341–73
  achieving 348–54
  benefits 363
  continuous improvement 360
  disadvantages 363–5
  effects 354–63
  employees 361–2
  extensions 365–6
  features 345–8
  kanbans 349–54
  principles 341–4
  pull methods 348–9
  total quality management 356–63
  suppliers 354–6
  view of stock 344

kaizen see continuous improvement
kanbans 349–54
  movement 351
  production 351
known demand 99–144
  constraints on average investment 132–4
  constraints on space 129–32

known demand (*continued*)
discounts in unit cost 100–12
discrete variable demand 135–41
finite replenishment rate 113–20
price discounts 100–12
shortages with back-orders 120–4
shortages with lost sales 125–8
variable reorder cost 109–12

last-in-first-out 49–51, 204
lead time 89–96, 149
causes of 89–90
constant 90–4
definition 89
demand 90–4, 149, 173–81
longer than stock cycle 92–4
Normally distributed 179–81
shorter than stock cycle 89–91
uncertain 179–80
variable 95
lead time demand 90–4, 149, 173–81
lean operations 40–1
length of supply chain 12–14
levels of decision 36–8, 268–70
LIFO 49–51, 204
Linear programming 280–1
linear regression 241–8
calculations 241–3
coefficient of correlation 245–7
coefficient of determination 244
extensions 247–8
purpose 241–2
spreadsheets 243, 247
location 33
logistics 11–19, 31–6
activities in 32–3
communications 33
integrating 34–6
role of inventory control 11–12
transport 32–3
warehousing 218–24
lost sales 120–1, 125–8
lumpy demand 164–5

MAD 237–8
management information system 196
manufacturing resource planning 332–4
marginal analysis 154–6
market survey 235–6
master schedule 274, 285–8
designing 286–7
spreadsheets 286–7
material, definition 5
material handling 32
material requirements planning 307–40
batching rules 327–9
benefits 321–4
bill of materials 311–21
capacity requirements planning 330
closed loop 330–1
definition 310–11
dependent demand methods 307–8
distribution resource planning 334
ERP 334–5
extensions of method 332–6
inputs and outputs 321–2
lot sizing 327–8
MRP II 332–4
problems 324–7
procedure 311–21
safety stock 323
spreadsheets 314, 316–18, 328–9
time shifting 312–13
mean absolute deviation 239–40
mean error 239
mean squared error 240
min-max system 186
MIS 196
mission 37
models for
dependent demand 305–73
economic order quantity 65–98
independent demand 63–190
known demand 99–144
constraints on average investment 132–4
constraints on space 129–32
discounts in unit cost 100–12
discrete variable demand 135–41
finite replenishment rate 113–20
price discounts 100–12
shortages with back-orders 120–4
shortages with lost sales 125–8
variable reorder cost 109–12
uncertain demand 147–89
discrete demand 154–70
intermittent demand 164–5
marginal analysis 154–6
Newsboy problem 156–61
periodic review 181–5
service level 170–2
shortages 166–70

single period problems 156–61
uncertain lead time demand 173–80
moving averages 249–52
MRP *see* material requirements planning
MRP II 332–4
multiple regression 247–8

national stocks 25–8
newsboy problem 156–61
spreadsheet for calculations 159–60
niche supplier 40
noise in forecasts 239
non-linear regression 248
Normally distributed
demand 152–3, 174–8, 180–1, 183–5
lead time 179–81

operational plans 268–70, 288–92
operational scheduling 268–70, 288–92
control 291–2
methods 288
scheduling rules 288–91
spreadsheets 290–1
order picking 33
order quantity
economic order quantity 65–98
with back-orders 120–8, 276
with discrete variable demand 135–41
with finite production rate 113–20
see also models for
organization, definition 5
outsourcing 21–2
ownership concentration 21

packaging 223–4
panel consensus 235
Pareto analysis 207–10
partnership 18
parts list 311–21
periodic review methods 181–5
advantages 185–6
spreadsheet 185
target stock level 181–5
with reorder level 186
personal insight for forecasting 235
pipeline stock 10
planned shortages 120–8
planning 267–302
aggregate plans 274–85
capacity 270–3
forecasts 260–2
levels 268–70
master schedule 285–8
operational schedules 288–92
strategic plans 268–70
tactical plans 274–88
postponement 23
price discounts 100–12
probabilistic models 86–7, 147–89
see also uncertain demand
procurement 32, 211–18
aims 211–12
cycle 213–15
definition 211
e-procurement 20, 216–17
procedure 213–15
supplier 213
product differentiation 40
production at finite rate 113–20
spreadsheet for calculations 117–19
projective forecasts 233–4, 248–60
exponential smoothing 252–5
moving averages 249–52
seasonality and trend 256–9
simple averages 248–9
tracking signal 254–5
pull methods 348–9
purchasing 32, 211–18
see also procurement

qualitative forecasts *see* judgemental forecasts
quantitative forecasts
causal 233–4, 241–8
projective 233–4, 248–60
time series 237–41
quick response 365–6

raw materials 9
receiving 32
recycling 33
reorder cost 52, 69
variable 109–12
reorder level 90–4
definition 90
gross stock 92–4
with periodic review method 186
replenishment 5–6
finite rate 113–20
resource requirements planning 333
return on assets 45–8
ROA 45–8
rule of 80–20 207–10

safety stock 9, 171–81, 183–5
definition 172
with material requirements planning 323
with uncertain demand 173–8
with uncertain lead time 179–81
schedules *see* operational scheduling
scheduling rules 288–91
scientific inventory control 24
seasonal index 256–9
seasonal stock 10
seasonality, forecasting 256–9
service level 164–5, 170–81
cycle 172–3
definition 171–2
sensitivity of
economic order quantity 78–82
forecasts 250, 254
shape of supply chain 12–14
short term scheduling *see* operational scheduling
shortage cost 53, 69, 122, 125, 161
shortages with
back-orders 120–4
uncertain demand 161–4, 166–70
lost sales 120–1, 125–8
simple averages for forecasting 248–9
simulation 292–9
approach 292–3
flow diagrams 293
packages 296–8
spreadsheets 294–5
single period model 156–61
spreadsheet for calculations 159–60
SKU *see* stock keeping unit
smoothing constant 253
spare parts 9, 161–5
spreadsheets for
aggregate planning 276–8, 282–4
discounts in unit costs 108–9
discrete variable demand 139–40
economic order quantity 75–7
errors in forecasts 240
finite production rate 117–19
linear regression 243, 247
master schedule 286–7
material requirements planning 314, 316–18, 328–9
newsboy problem 159–60
Normally distributed demand 178
operational schedules 290–1
periodic review method 185
projective forecasting 251, 254–5, 257–8
simulation 294–5
single period model 159–60
TQM costs 358–9
varying demand 15–17
varying order size 83–4
stochastic models 86–7, 147–87
see also uncertain demand
stock
aggregate 24–8
as buffer 7–9
classification 9–10
co-managed 22
control *see* inventory management
costs of holding 48–53
cycle 5–7, 66, 70
definition 3–4
gross 92–4
JIT view 344
national 25–8
reasons for holding 7–9
replenishment 5–6
strategic role 44–8
trends affecting 19–24
types of 9–10
value 48–51
views of 24, 344
stock control *see* inventory management
stock cycle 5–7, 66, 70
stock keeping unit 4
stock/non-stock decision 55–6, 125–8
stocktaking 200–1
storage space constraint 129–32
strategic alliance 18
strategic decisions 37, 268–70
strategic role of stock 44–8
strategy
agile 41, 268–9
alternative 39–44
business 37–44
corporate 37–44
cost leadership 40
focus 41–2, 268–9
functional 37
implementing 42
lean 40–1, 268–9
niche supplier 40
product differentiation 40
supplier
choosing 213
definition 7
relations with JIT 354–6
supply and demand 211–17
supply chain
benefits of co-operation 17–18
breadth 12–14
co-operation in 15–19
demand amplification 15–17

design 13–15
length 12–14
management 11–19, 31–6
problems with co-operation 15–17
shape of 12–14
supply chain management 11–19, 31–6
see also logistics

tactical plans 268–70, 274
Taguchi 148
target stock level 181–5
third party operators 22
three-bin system 95, 178
time horizons of forecasts 233
time series 237–41
time shifting in MRP 312–13
total quality management 356–63
achieving 359–63
costs 357
jidoka 360
spreadsheet for costs 358–9
TQM *see* total quality management
tracking signal 254–5
trading accounts 202–3
transaction recording 198–201
transport 32–3
trends
affecting stock 19–24
in time series 256–9
two-bin system 95, 178
types of
cost 52–3
customer 9–10
forecasting method 233–4
inventory management method 57–8
stock 9–10

uncertain demand 86–7, 147–91
discrete demand 154–70
intermittent demand 164–5
marginal analysis 154–6
Newsboy problem 156–61
periodic review 181–5
service level 170–2
shortages 166–70
single period problems 156–61
uncertain lead time demand 173–80
uncertainty
both lead time and demand 180–1
causes 148–50
costs 87
demand *see* uncertain demand
lead time 179–80
lead time demand 173–81
reducing 148–9
unit 4
unit cost 52, 69
discounts 100–12
unitization 223
utility 43

valid cost curve 100–12
valid minimum 102–12
value of stock 48–51, 204–7
inferred 206–7
variable
cost 71–3, 100
demand 15–17
lead time 95
unit cost 100–12
varying order size 83–4
vendor managed inventory 22
vision 37
VMI 22

warehousing 32, 218–24
activities 218–19
automation 222–3
cross docking 22–3, 219
design 220–3
packaging 223–4
purpose 218
unitization 223
waste 40–1
weaknesses of
co-operation in supply chain 15–19
economic order quantity 77–8
independent demand methods 307–8
just-in-time 363–5
MRP 324–7
weighted average cost of stock 49–51, 205
work in process 9
work in progress 9

VELVSHER